



US00934447B2

(12) **United States Patent**
Cohen et al.

(10) **Patent No.:** **US 9,344,447 B2**
(45) **Date of Patent:** ***May 17, 2016**

(54) **INTERNAL MALWARE DATA ITEM
CLUSTERING AND ANALYSIS**

(71) Applicant: **Palantir Technologies Inc.**, Palo Alto,
CA (US)

(72) Inventors: **David Cohen**, Mountain View, CA (US);
Jason Ma, Mountain View, CA (US);
Bing Jie Fu, Redwood City, CA (US);
Ilya Nepomnyashchii, Mountain View,
CA (US); **Steven Berler**, Menlo Park,
CA (US); **Alex Smaliy**, Palo Alto, CA
(US); **Jack Grossman**, San Francisco,
CA (US); **James Thompson**, San
Francisco, CA (US); **Julia Boortz**,
Menlo Park, CA (US); **Matthew
Sprague**, Palo Alto, CA (US); **Parvathy
Menon**, Palo Alto, CA (US); **Michael
Kross**, Palo Alto, CA (US); **Michael
Harris**, Palo Alto, CA (US); **Adam
Borochoff**, New York, NY (US)

(73) Assignee: **Palantir Technologies Inc.**, Palo Alto,
CA (US)

(*) Notice: Subject to any disclaimer, the term of this
patent is extended or adjusted under 35
U.S.C. 154(b) by 161 days.

This patent is subject to a terminal dis-
claimer.

(21) Appl. No.: **14/486,991**

(22) Filed: **Sep. 15, 2014**

(65) **Prior Publication Data**

US 2016/0006749 A1 Jan. 7, 2016

Related U.S. Application Data

(63) Continuation of application No. 14/473,920, filed on
Aug. 29, 2014.

(60) Provisional application No. 62/020,876, filed on Jul. 3,
2014.

(51) **Int. Cl.**

G06Q 40/00 (2012.01)

H04L 29/06 (2006.01)

G06F 17/30 (2006.01)

(52) **U.S. Cl.**

CPC **H04L 63/145** (2013.01); **G06F 17/30598**
(2013.01); **H04L 63/1408** (2013.01)

(58) **Field of Classification Search**

USPC 705/35, 36 R, 37, 38
See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

5,109,399 A 4/1992 Thompson
5,329,108 A 7/1994 Lamoure

(Continued)

FOREIGN PATENT DOCUMENTS

EP 1 191 463 3/2002
EP 1 672 527 6/2006

(Continued)

OTHER PUBLICATIONS

"A First Look: Predicting Market Demand for Food Retail using a
Huff Analysis," TRF Policy Solutions, Jul. 2012, pp. 30.

(Continued)

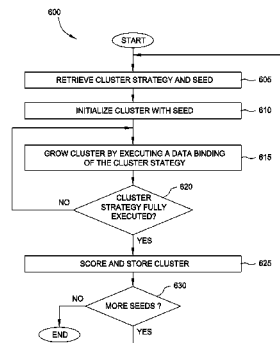
Primary Examiner — Lalita M Hamilton

(74) *Attorney, Agent, or Firm* — Knobbe, Martens, Olson &
Bear, LLP

(57) **ABSTRACT**

Embodiments of the present disclosure relate to a data analy-
sis system that may automatically generate memory-efficient
clustered data structures, automatically analyze those clus-
tered data structures, and provide results of the automated
analysis in an optimized way to an analyst. The automated
analysis of the clustered data structures (also referred to
herein as data clusters) may include an automated application
of various criteria or rules so as to generate a compact,
human-readable analysis of the data clusters. The human-
readable analysis (also referred to herein as "summaries" or
"conclusions") of the data clusters may be organized into an
interactive user interface so as to enable an analyst to quickly
navigate among information associated with various data
clusters and efficiently evaluate those data clusters in the
context of, for example, a fraud investigation. Embodiments
of the present disclosure also relate to automated scoring of
the clustered data structures.

19 Claims, 37 Drawing Sheets



(56)

References Cited

U.S. PATENT DOCUMENTS

5,632,009 A	5/1997	Rao et al.	8,135,679 B2	3/2012	Bayliss
5,670,987 A	9/1997	Doi et al.	8,135,719 B2	3/2012	Bayliss
5,845,300 A	12/1998	Comer et al.	8,145,703 B2	3/2012	Frishert et al.
6,057,757 A	5/2000	Arrowsmith et al.	8,214,361 B1	7/2012	Sandler et al.
6,091,956 A	7/2000	Hollenberg	8,214,764 B2	7/2012	Gemmell et al.
6,161,098 A	12/2000	Wallman	8,225,201 B2	7/2012	Michael
6,219,053 B1	4/2001	Tachibana et al.	8,229,947 B2	7/2012	Fujinaga
6,232,971 B1	5/2001	Haynes	8,230,333 B2	7/2012	Decherd et al.
6,247,019 B1	6/2001	Davies	8,266,168 B2	9/2012	Bayliss
6,279,018 B1	8/2001	Kudrolli et al.	8,280,880 B1	10/2012	Aymeloglu et al.
6,341,310 B1	1/2002	Leshem et al.	8,290,942 B2	10/2012	Jones et al.
6,369,835 B1	4/2002	Lin	8,301,464 B1	10/2012	Cave et al.
6,549,944 B1	4/2003	Weinberg et al.	8,321,943 B1	11/2012	Walters et al.
6,567,936 B1	5/2003	Yang et al.	8,347,398 B1	1/2013	Weber
6,581,068 B1	6/2003	Bensoussan et al.	8,352,881 B2	1/2013	Champion et al.
6,631,496 B1	10/2003	Li et al.	8,368,695 B2	2/2013	Howell et al.
6,714,936 B1	3/2004	Nevin, III	8,397,171 B2	3/2013	Klassen et al.
6,775,675 B1	8/2004	Nwabueze et al.	8,412,707 B1	4/2013	Mianji
6,839,745 B1	1/2005	Dingari et al.	8,447,674 B2	5/2013	Choudhuri et al.
6,877,137 B1	4/2005	Rivette et al.	8,447,722 B1	5/2013	Ahuja et al.
6,976,210 B1	12/2005	Silva et al.	8,452,790 B1	5/2013	Mianji
6,985,950 B1	1/2006	Hanson et al.	8,463,036 B1	6/2013	Ramesh et al.
7,036,085 B2	4/2006	Barros	8,484,168 B2	7/2013	Bayliss
7,043,702 B2	5/2006	Chi et al.	8,489,331 B2	7/2013	Kopf et al.
7,055,110 B2	5/2006	Kupka et al.	8,489,641 B1	7/2013	Seefeld et al.
7,139,800 B2	11/2006	Bellotti et al.	8,495,077 B2	7/2013	Bayliss
7,158,878 B2	1/2007	Rasmussen et al.	8,498,969 B2	7/2013	Bayliss
7,162,475 B2	1/2007	Ackerman	8,498,984 B1	7/2013	Hwang et al.
7,168,039 B2	1/2007	Bertram	8,514,082 B2	8/2013	Cova et al.
7,171,427 B2	1/2007	Witkowski et al.	8,515,207 B2	8/2013	Chau
7,269,786 B1	9/2007	Malloy et al.	8,554,579 B2	10/2013	Tribble et al.
7,278,105 B1	10/2007	Kitts	8,577,911 B1	11/2013	Stepinski et al.
7,290,698 B2	11/2007	Poslinski et al.	8,589,273 B2	11/2013	Creeden et al.
7,333,998 B2	2/2008	Heckerman et al.	8,600,872 B1	12/2013	Yan
7,373,669 B2	5/2008	Eisen	8,620,641 B2	12/2013	Farnsworth et al.
7,379,811 B2	5/2008	Rasmussen et al.	8,689,108 B1	4/2014	Duffield et al.
7,379,903 B2	5/2008	Caballero et al.	8,713,467 B1	4/2014	Goldenberg et al.
7,426,654 B2	9/2008	Adams et al.	8,739,278 B2	5/2014	Varghese
7,451,397 B2	11/2008	Weber et al.	8,745,516 B2	6/2014	Mason et al.
7,454,466 B2	11/2008	Bellotti et al.	8,781,169 B2	7/2014	Jackson et al.
7,467,375 B2	12/2008	Tondreau et al.	8,788,405 B1	7/2014	Sprague et al.
7,502,786 B2	3/2009	Liu et al.	8,788,407 B1	7/2014	Singh et al.
7,525,422 B2	4/2009	Bishop et al.	8,799,799 B1	8/2014	Cervelli et al.
7,529,727 B2	5/2009	Arning et al.	8,812,960 B1	8/2014	Sun et al.
7,558,677 B2	7/2009	Jones	8,818,892 B1	8/2014	Sprague et al.
7,574,428 B2	8/2009	Leiserowitz et al.	8,832,594 B1	9/2014	Thompson et al.
7,617,232 B2	11/2009	Gabbert et al.	2001/0021936 A1	9/2001	Bertram
7,620,628 B2	11/2009	Kapur et al.	2002/0091707 A1	7/2002	Keller
7,627,812 B2	12/2009	Chamberlain et al.	2002/0095658 A1	7/2002	Shulman
7,634,717 B2	12/2009	Chamberlain et al.	2002/0116120 A1	8/2002	Ruiz et al.
7,703,021 B1	4/2010	Flam	2002/0130907 A1	9/2002	Chi et al.
7,712,049 B2	5/2010	Williams et al.	2002/0174201 A1	11/2002	Ramer et al.
7,716,077 B1	5/2010	Mikurak	2003/0039948 A1	2/2003	Donahue
7,725,547 B2	5/2010	Albertson et al.	2003/0097330 A1	5/2003	Hillmer et al.
7,770,100 B2	8/2010	Chamberlain et al.	2003/0144868 A1	7/2003	MacIntyre et al.
7,783,658 B1	8/2010	Bayliss	2003/0163352 A1	8/2003	Surpin et al.
7,809,703 B2	10/2010	Balabhadrapatruni et al.	2003/0200217 A1	10/2003	Ackerman
7,814,102 B2	10/2010	Miller et al.	2003/0225755 A1	12/2003	Iwayama et al.
7,818,658 B2	10/2010	Chen	2004/0064256 A1	4/2004	Barinek et al.
7,894,984 B2	2/2011	Rasmussen et al.	2004/0085318 A1	5/2004	Hassler et al.
7,899,611 B2	3/2011	Downs et al.	2004/0095349 A1	5/2004	Bito et al.
7,917,376 B2	3/2011	Bellin et al.	2004/0143602 A1	7/2004	Ruiz et al.
7,920,963 B2	4/2011	Jouline et al.	2004/0181554 A1	9/2004	Heckerman et al.
7,933,862 B2	4/2011	Chamberlain et al.	2004/0205524 A1	10/2004	Richter et al.
7,962,281 B2	6/2011	Rasmussen et al.	2004/0260702 A1	12/2004	Cragun et al.
7,962,848 B2	6/2011	Bertram	2005/0027705 A1	2/2005	Sadri et al.
8,001,465 B2	8/2011	Kudrolli et al.	2005/0028094 A1	2/2005	Allyn
8,001,482 B2	8/2011	Bhattiprolu et al.	2005/0080769 A1	4/2005	Gemmell
8,010,545 B2	8/2011	Stefik et al.	2005/0086207 A1	4/2005	Heuer et al.
8,015,487 B2	9/2011	Roy et al.	2005/0125715 A1	6/2005	Franco et al.
8,024,778 B2	9/2011	Cash et al.	2005/0180330 A1	8/2005	Shapiro
8,036,632 B1	10/2011	Cona et al.	2005/0182793 A1	8/2005	Keenan et al.
8,046,362 B2	10/2011	Bayliss	2005/0183005 A1	8/2005	Denoue et al.
8,103,543 B1	1/2012	Zwicky	2005/0251786 A1	11/2005	Citron et al.
8,134,457 B2	3/2012	Velipasalar et al.	2006/0026170 A1	2/2006	Kreitler et al.
			2006/0045470 A1	3/2006	Poslinski et al.
			2006/0059139 A1	3/2006	Robinson
			2006/0074866 A1	4/2006	Chamberlain et al.
			2006/0080619 A1	4/2006	Carlson et al.

(56)

References Cited

U.S. PATENT DOCUMENTS

2006/0129746	A1	6/2006	Porter	2011/0119100	A1	5/2011	Ruhl et al.
2006/0139375	A1	6/2006	Rasmussen et al.	2011/0137766	A1	6/2011	Rasmussen et al.
2006/0149596	A1	7/2006	Surpin et al.	2011/0167105	A1	7/2011	Ramakrishnan et al.
2006/0203337	A1	9/2006	White	2011/0167493	A1	7/2011	Song et al.
2006/0241974	A1	10/2006	Chao et al.	2011/0170799	A1	7/2011	Carrino et al.
2006/0271277	A1	11/2006	Hu et al.	2011/0173093	A1	7/2011	Psota et al.
2007/0011150	A1	1/2007	Frank	2011/0178842	A1	7/2011	Rane et al.
2007/0016363	A1	1/2007	Huang et al.	2011/0208724	A1	8/2011	Jones et al.
2007/0038962	A1	2/2007	Fuchs et al.	2011/0231223	A1	9/2011	Winters
2007/0078832	A1	4/2007	Ott et al.	2011/0238510	A1	9/2011	Rowen et al.
2007/0192265	A1	8/2007	Chopin et al.	2011/0238570	A1	9/2011	Li et al.
2007/0208497	A1	9/2007	Downs et al.	2011/0291851	A1	12/2011	Whisenant
2007/0208498	A1	9/2007	Barker et al.	2011/0307382	A1	12/2011	Siegel et al.
2007/0266336	A1	11/2007	Nojima et al.	2011/0310005	A1	12/2011	Chen et al.
2007/0294200	A1	12/2007	Au	2012/0019559	A1	1/2012	Siler et al.
2007/0294643	A1	12/2007	Kyle	2012/0036434	A1	2/2012	Oberstein
2008/0040684	A1	2/2008	Crump	2012/0066296	A1	3/2012	Appleton et al.
2008/0051989	A1	2/2008	Welsh	2012/0079363	A1	3/2012	Folting et al.
2008/0052142	A1	2/2008	Bailey et al.	2012/0084135	A1	4/2012	Nissan et al.
2008/0077597	A1	3/2008	Butler	2012/0084866	A1	4/2012	Stolfo
2008/0077642	A1	3/2008	Carbone et al.	2012/0106801	A1	5/2012	Jackson
2008/0104019	A1	5/2008	Nath	2012/0144335	A1	6/2012	Abeln et al.
2008/0126951	A1	5/2008	Sood et al.	2012/0159307	A1	6/2012	Chung et al.
2008/0133567	A1	6/2008	Ames et al.	2012/0196557	A1	8/2012	Reich et al.
2008/0195417	A1	8/2008	Surpin et al.	2012/0196558	A1	8/2012	Reich et al.
2008/0195608	A1	8/2008	Clover	2012/0208636	A1	8/2012	Feige
2008/0222295	A1	9/2008	Robinson et al.	2012/0215898	A1	8/2012	Shah et al.
2008/0263468	A1	10/2008	Cappione et al.	2012/0221511	A1	8/2012	Gibson et al.
2008/0267107	A1	10/2008	Rosenberg	2012/0221580	A1	8/2012	Barney
2008/0278311	A1	11/2008	Grange et al.	2012/0246148	A1	9/2012	Dror
2008/0288306	A1	11/2008	MacIntyre et al.	2012/0290879	A1	11/2012	Shibuya et al.
2008/0301643	A1	12/2008	Appleton et al.	2012/0310831	A1	12/2012	Harris et al.
2009/0002492	A1	1/2009	Velipasalar et al.	2012/0310838	A1	12/2012	Harris et al.
2009/0018940	A1	1/2009	Wang et al.	2012/0323888	A1	12/2012	Osann, Jr.
2009/0027418	A1	1/2009	Maru et al.	2012/0330973	A1	12/2012	Ghuneim et al.
2009/0030915	A1	1/2009	Winter et al.	2013/0006725	A1	1/2013	Simanek et al.
2009/0044279	A1	2/2009	Crawford et al.	2013/0024307	A1	1/2013	Fuerstenberg et al.
2009/0076845	A1	3/2009	Bellin et al.	2013/0024339	A1	1/2013	Choudhuri et al.
2009/0083184	A1	3/2009	Eisen	2013/0046842	A1	2/2013	Muntz et al.
2009/0119309	A1	5/2009	Gibson et al.	2013/0061169	A1	3/2013	Pearcy et al.
2009/0125369	A1	5/2009	Klooststra et al.	2013/0073377	A1	3/2013	Heath
2009/0132953	A1	5/2009	Reed et al.	2013/0078943	A1	3/2013	Biage et al.
2009/0144262	A1	6/2009	White et al.	2013/0097482	A1	4/2013	Marantz et al.
2009/0164934	A1	6/2009	Bhattiprolu et al.	2013/0101159	A1	4/2013	Chao et al.
2009/0171939	A1	7/2009	Athsani et al.	2013/0150004	A1	6/2013	Rosen
2009/0192957	A1	7/2009	Subramanian et al.	2013/0160120	A1	6/2013	Malaviya et al.
2009/0222400	A1	9/2009	Kupersmidt et al.	2013/0176321	A1	7/2013	Mitchell et al.
2009/0234720	A1	9/2009	George et al.	2013/0179420	A1	7/2013	Park et al.
2009/0281839	A1	11/2009	Lynn et al.	2013/0211985	A1	8/2013	Clark et al.
2009/0287470	A1	11/2009	Farnsworth et al.	2013/0224696	A1	8/2013	Wolfe et al.
2009/0292626	A1	11/2009	Oxford	2013/0232045	A1	9/2013	Tai et al.
2010/0011282	A1	1/2010	Dollard et al.	2013/0238616	A1	9/2013	Rose et al.
2010/0042922	A1	2/2010	Bradeteanu et al.	2013/0262527	A1	10/2013	Hunter et al.
2010/0057716	A1	3/2010	Stefik et al.	2013/0268520	A1	10/2013	Fisher et al.
2010/0070523	A1	3/2010	Delgo et al.	2013/0282696	A1	10/2013	John et al.
2010/0070897	A1	3/2010	Aymeloglu et al.	2013/0290011	A1	10/2013	Lynn et al.
2010/0077483	A1	3/2010	Stolfo et al.	2013/0290825	A1	10/2013	Arndt et al.
2010/0122152	A1	5/2010	Chamberlain et al.	2013/0297619	A1	11/2013	Chandrasekaran et al.
2010/0125546	A1	5/2010	Barrett et al.	2013/0318594	A1	11/2013	Hoy et al.
2010/0131457	A1	5/2010	Heimendinger	2013/0339218	A1	12/2013	Subramanian et al.
2010/0185691	A1	7/2010	Irmak et al.	2014/0006109	A1	1/2014	Callioni et al.
2010/0198684	A1	8/2010	Eraker et al.	2014/0019936	A1	1/2014	Cohanoff
2010/0199225	A1	8/2010	Coleman et al.	2014/0032506	A1	1/2014	Hoey et al.
2010/0250412	A1	9/2010	Wagner	2014/0033010	A1	1/2014	Richardt et al.
2010/0280857	A1	11/2010	Liu et al.	2014/0047357	A1	2/2014	Alfaro et al.
2010/0293174	A1	11/2010	Bennett et al.	2014/0068487	A1	3/2014	Steiger et al.
2010/0306029	A1	12/2010	Jolley	2014/0108380	A1	4/2014	Gotz et al.
2010/0306713	A1	12/2010	Geisner et al.	2014/0108985	A1	4/2014	Scott et al.
2010/0321399	A1	12/2010	Ellren et al.	2014/0157172	A1	6/2014	Peery et al.
2010/0325581	A1	12/2010	Finkelstein et al.	2014/0195515	A1	7/2014	Baker et al.
2011/0047159	A1	2/2011	Baid et al.				
2011/0060753	A1	3/2011	Shaked et al.				
2011/0061013	A1	3/2011	Billicki et al.				
2011/0074811	A1	3/2011	Hanson et al.				
2011/0087519	A1	4/2011	Fordyce, III et al.				

FOREIGN PATENT DOCUMENTS

EP	2 551 799	1/2013
WO	WO 2005/104736	11/2005
WO	WO 2008/011728	1/2008
WO	WO 2009/061501	5/2009

(56)

References Cited

FOREIGN PATENT DOCUMENTS

WO	WO 2010/000014	1/2010
WO	WO 2010/030913	3/2010
WO	WO 2013/126281	8/2013

OTHER PUBLICATIONS

"A Quick Guide to UniProtKB Swiss-Prot & TrEMBL," Sep. 2011, pp. 2.

Acklen, Laura, "Absolute Beginner's Guide to Microsoft Word 2003," Dec. 24, 2003, pp. 15-18, 34-41, 308-316.

Alfred, Rayner "Summarizing Relational Data Using Semi-Supervised Genetic Algorithm-Based Clustering Techniques", Journal of Computer Science, 2010, vol. 6, No. 7, pp. 775-784.

Ananiev et al., "The New Modality API," <http://web.archive.org/web/20061211011958/http://java.sun.com/developer/technicalArticles/J2SE/Desktop/javase6/modality/> Jan. 21, 2006, pp. 8.

Bluttmann et al., "Excel Formulas and Functions for Dummies," 2005, Wiley Publishing, Inc., pp. 280, 284-286.

Bugzilla@Mozilla, "Bug 18726—[feature] Long-click means of invoking contextual menus not supported," http://bugzilla.mozilla.org/show_bug.cgi?id=18726 printed Jun. 13, 2013 in 11 pages.

Canese et al., "Chapter 2: PubMed: The Bibliographic Database," The NCBI Handbook, Oct. 2002, pp. 1-10.

Chen et al., "Bringing Order to the Web: Automatically Categorizing Search Results," CHI 2000, Proceedings of the SIGCHI conference on Human Factors in Computing Systems, Apr. 1-6, 2000, The Hague, The Netherlands, pp. 145-152.

Conner, Nancy, "Google Apps: The Missing Manual," Sharing and Collaborating on Documents, May 1, 2008, pp. 93-97, 106-113 & 120-121.

Delcher et al., "Identifying Bacterial Genes and Endosymbiont DNA with Glimmer," Bioinformatics, vol. 23, No. 6, 2007, pp. 673-679.

Dramowicz, Ela, "Retail Trade Area Analysis Using the Huff Model," Directions Magazine, Jul. 2, 2005 in 10 pages, <http://www.directionsmag.com/articles/retail-trade-area-analysis-using-the-huff-model/123411>.

GIS-NET 3 Public—Department of Regional Planning. Planning & Zoning Information for Unincorporated LA County. Retrieved Oct. 2, 2013 from http://gis.planning.lacounty.gov/GIS-NET3_Public/Viewer.html.

Goswami, Gautam, "Quite 'Writely' Said!" One Brick at a Time, Aug. 21, 2005, pp. 7.

Griffith, Daniel A., "A Generalized Huff Model," Geographical Analysis, Apr. 1982, vol. 14, No. 2, pp. 135-144.

Hansen et al., "Analyzing Social Media Networks with NodeXL: Insights from a Connected World", Elsevier Science, Sep. 2010, Ch. 4 & 10, pp. 53-67 & 143-164.

Hibbert et al., "Prediction of Shopping Behavior Using a Huff Model Within a GIS Framework," Healthy Eating in Context, Mar. 18, 2011, pp. 16.

Huff et al., "Calibrating the Huff Model Using ArcGIS Business Analyst," ESRI, Sep. 2008, pp. 33.

Huff, David L., "Parameter Estimation in the Huff Model," ESRI, ArcUser, Oct.-Dec. 2003, pp. 34-36.

Kahan et al., "Annotea: An Open RDF Infrastructure for Shared WEB Annotations", Computer Networks, 2002, vol. 39, pp. 589-608.

Keylines.com, "An Introduction to KeyLines and Network Visualization," Mar. 2014, <http://keylines.com/wp-content/uploads/2014/03/KeyLines-White-Paper.pdf> downloaded May 12, 2014 in 8 pages.

Keylines.com, "KeyLines Datasheet," Mar. 2014, <http://keylines.com/wp-content/uploads/2014/03/KeyLines-datasheet.pdf> downloaded May 12, 2014 in 2 pages.

Keylines.com, "Visualizing Threats: Improved Cyber Security Through Network Visualization," Apr. 2014, <http://keylines.com/wp-content/uploads/2014/04/Visualizing-Threats1.pdf> downloaded May 12, 2014 in 10 pages.

Kitts, Paul, "Chapter 14: Genome Assembly and Annotation Process," The NCBI Handbook, Oct. 2002, pp. 1-21.

Liu, Tianshun, "Combining GIS and the Huff Model to Analyze Suitable Locations for a New Asian Supermarket in the Minneapolis and St. Paul, Minnesota USA," Papers in Resource Analysis, 2012, vol. 14, pp. 8.

Madden, Tom, "Chapter 16: The BLAST Sequence Analysis Tool," The NCBI Handbook, Oct. 2002, pp. 1-15.

Manno et al., "Introducing Collaboration in Single-user Applications through the Centralized Control Architecture," 2010, pp. 10.

Manske, "File Saving Dialogs," http://www.mozila.org/editor/ui_specs/FileSaveDialogs.html, Jan. 20, 1999, pp. 7.

Map of San Jose, CA. Retrieved Oct. 2, 2013 from <http://maps.bing.com>.

Map of San Jose, CA. Retrieved Oct. 2, 2013 from <http://maps.google.com>.

Map of San Jose, CA. Retrieved Oct. 2, 2013 from <http://maps.yahoo.com>.

Microsoft—Developer Network, "Getting Started with VBA in Word 2010," Apr. 2010, <http://msdn.microsoft.com/en-us/library/ff604039%28v=office.14%29.aspx> as printed Apr. 4, 2014 in 17 pages.

Microsoft Office—Visio, "About connecting shapes," <http://office.microsoft.com/en-us/visio-help/about-connecting-shapes-HP085050369.aspx> printed Aug. 4, 2011 in 6 pages.

Microsoft Office—Visio, "Add and glue connectors with the Connector tool," <http://office.microsoft.com/en-us/visio-help/add-and-glue-connectors-with-the-connector-tool-HA010048532.aspx?CTT=1> printed Aug. 4, 2011 in 1 page.

Mizrahi, Ilene, "Chapter 1: GenBank: The Nucleotide Sequence Database," The NCBI Handbook, Oct. 2002, pp. 1-14.

Rouse, Margaret, "OLAP Cube," <http://searchdatamanagement.techtarget.com/definition/OLAP-cube>, Apr. 28, 2012, pp. 16.

Shah, Chintan, "Periodic Connections to Control Server Offer New Way to Detect Botnets," Oct. 24, 2013 in 6 pages, <http://www.blogs.mcafee.com/mcafee-labs/periodic-links-to-control-server-offer-new-way-to-detect-botnets>.

Sigrist, et al., "PROSITE, a Protein Domain Database for Functional Characterization and Annotation," Nucleic Acids Research, 2010, vol. 38, pp. D161-D166.

Sirotkin et al., "Chapter 13: The Processing of Biological Sequence Data at NCBI," The NCBI Handbook, Oct. 2002, pp. 1-11.

"The FASTA Program Package," fasta-36.3.4, Mar. 25, 2011, pp. 29.

Wiggerts, T.A., "Using Clustering Algorithms in Legacy Systems Remodularization," Reverse Engineering, Proceedings of the Fourth Working Conference, Netherlands, Oct. 6-8, 1997, IEEE Computer Soc., pp. 33-43.

International Search Report and Written Opinion in Application No. PCT/US2009/056703, dated Mar. 15, 2010.

Official Communication in European Application No. EP 14158861.6 dated Jun. 16, 2014.

Official Communication in New Zealand Application No. 622517 dated Apr. 3, 2014.

Official Communication in New Zealand Application No. 624557 dated May 14, 2014.

Official Communication in New Zealand Application No. 628263 dated Aug. 12, 2014.

European Search Report in Application No. 14159535.5, dated May 22, 2014.

Official Communication in New Zealand Application No. 622181 dated Mar. 24, 2014.

Official Communication in New Zealand Application No. 627962 dated Aug. 5, 2014.




FIG. 1

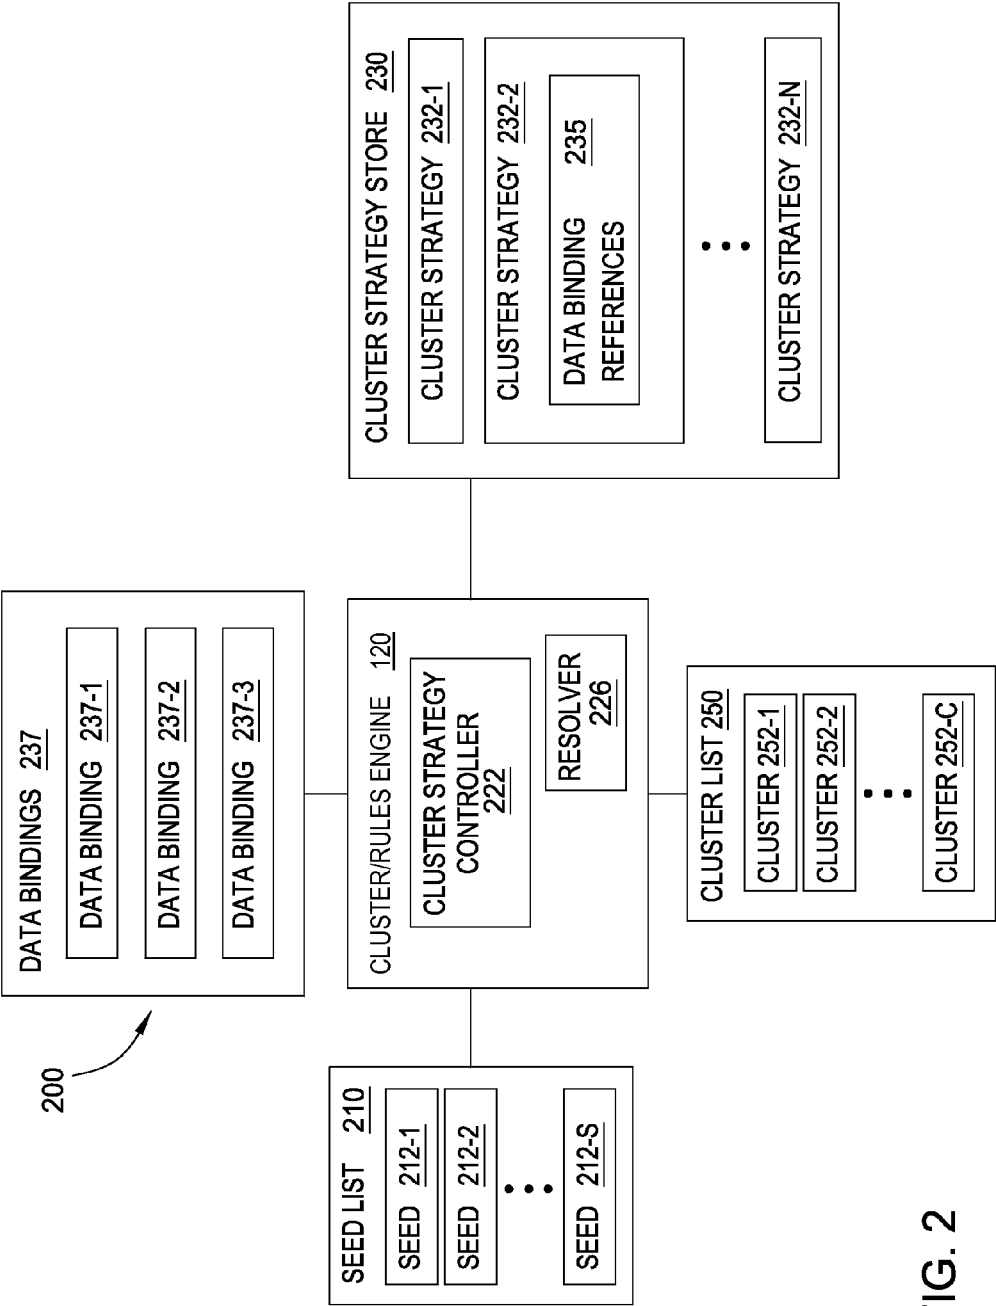


FIG. 2




FIG. 3A





FIG. 3B



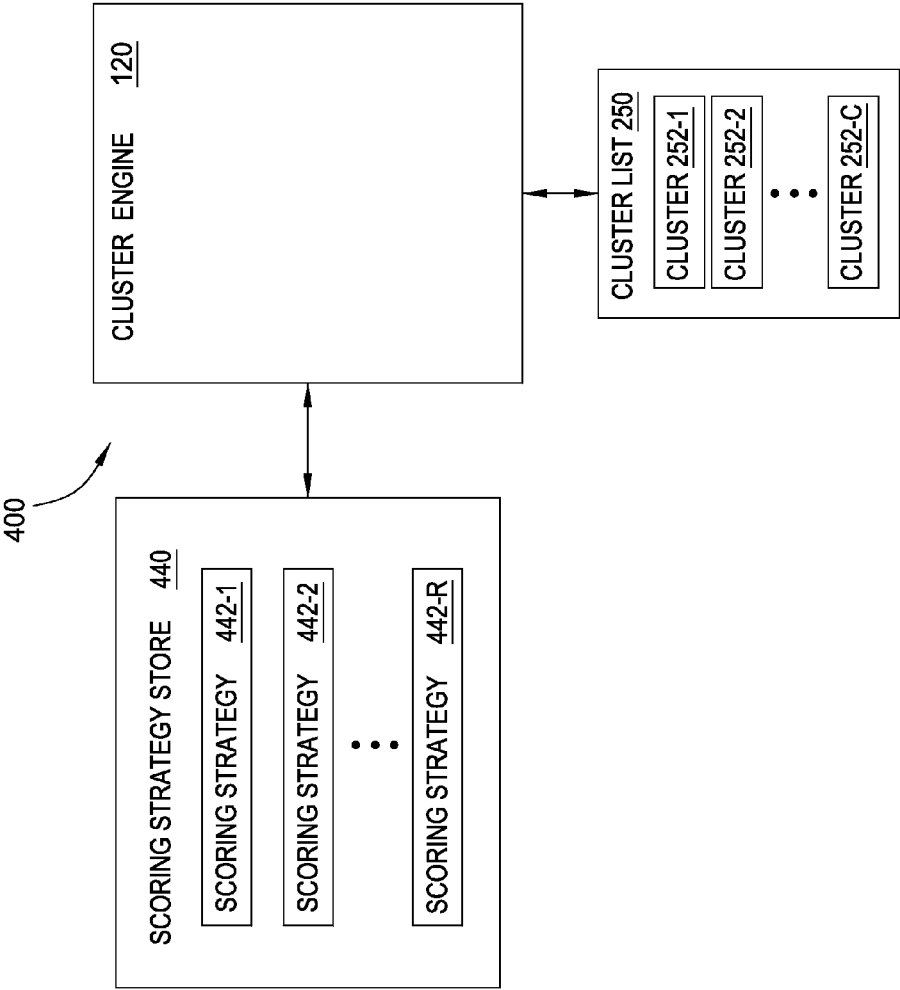


FIG. 4




FIG. 5

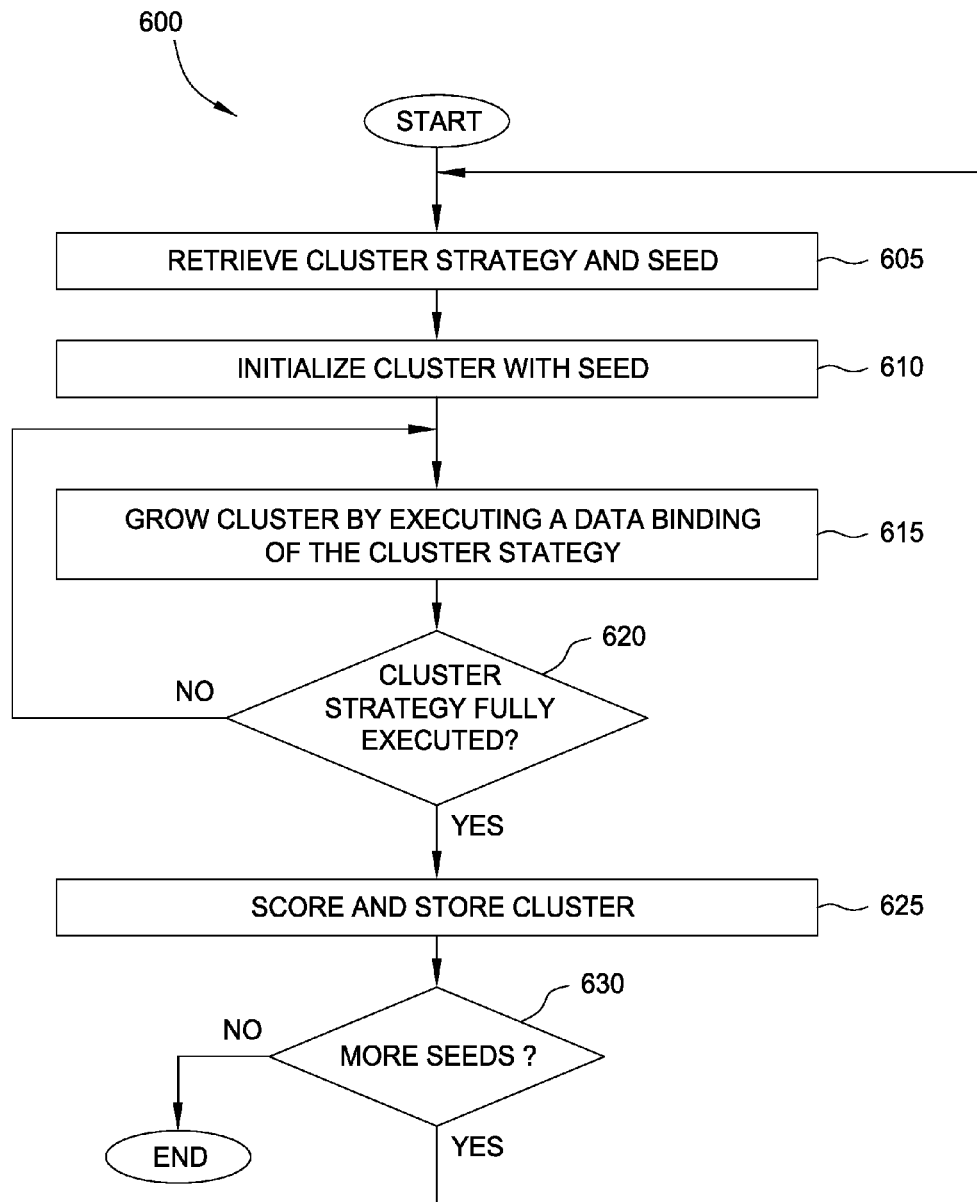


FIG. 6

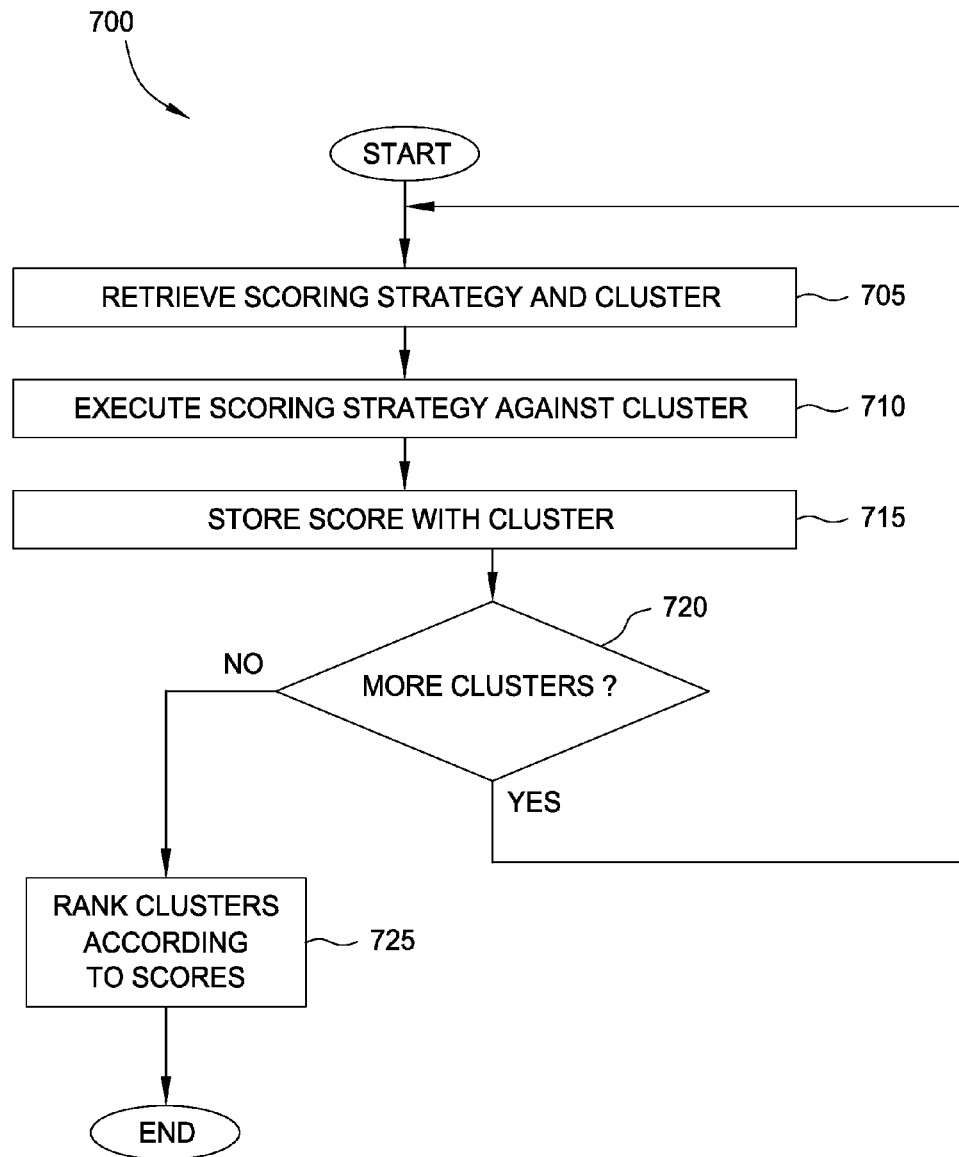
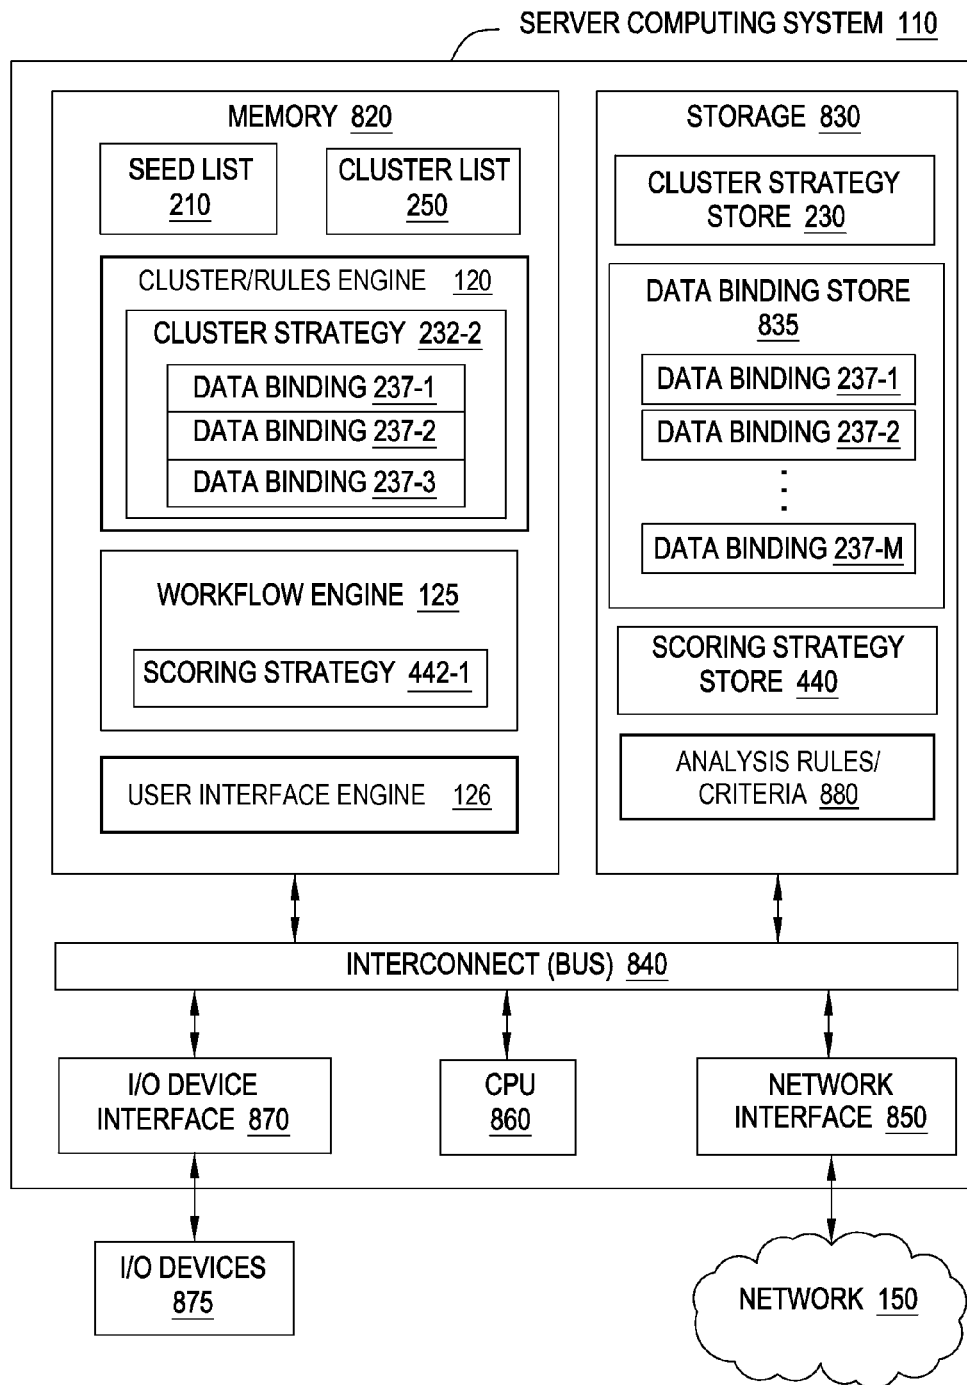
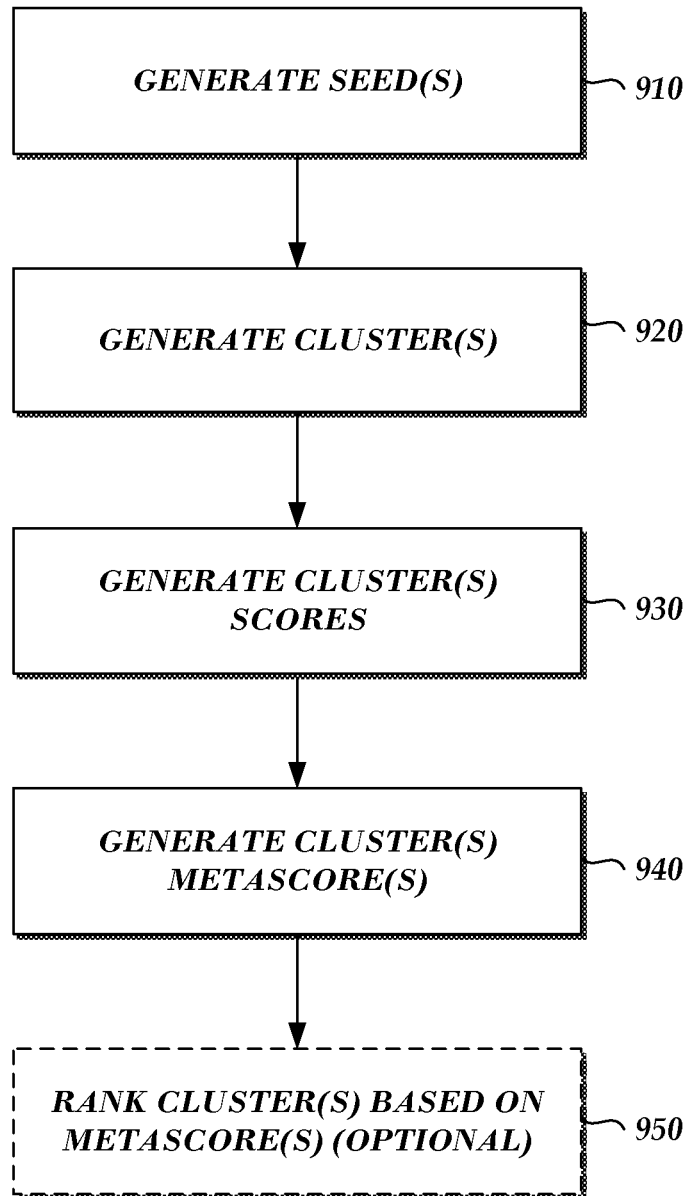


FIG. 7



**FIG. 9**

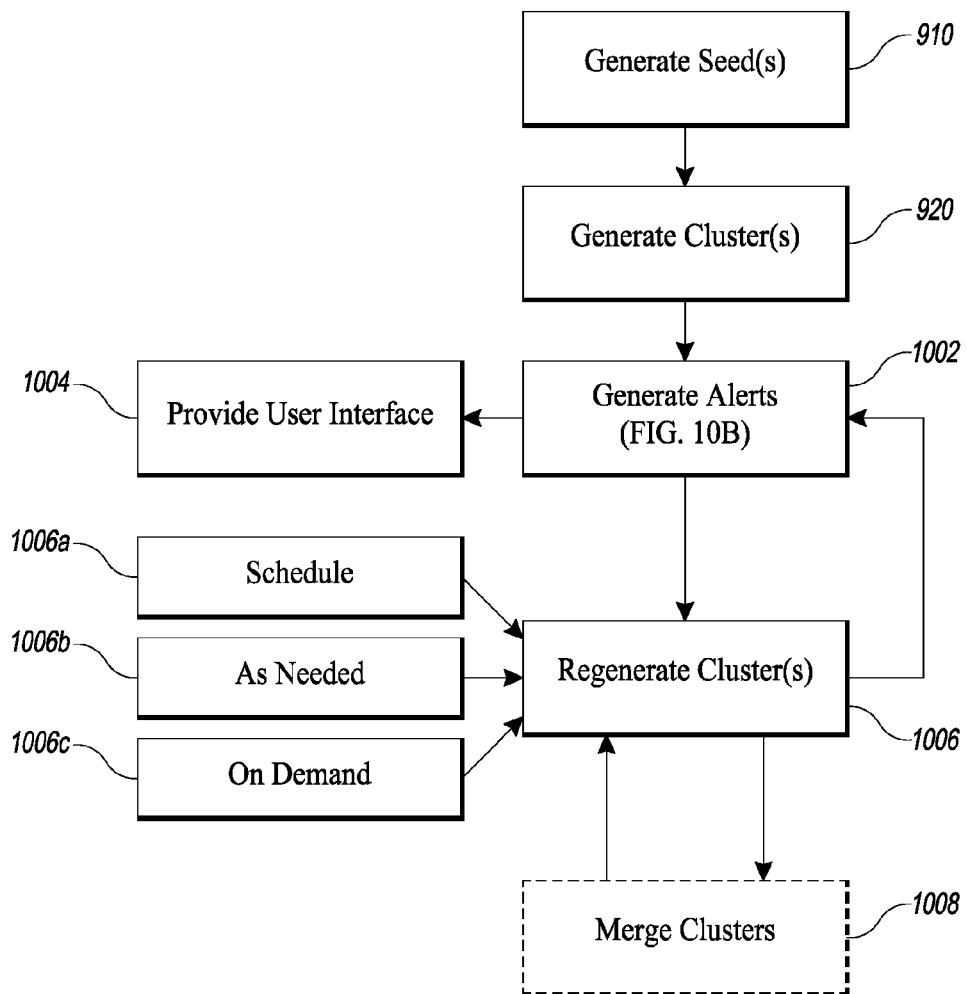


FIG. 10A

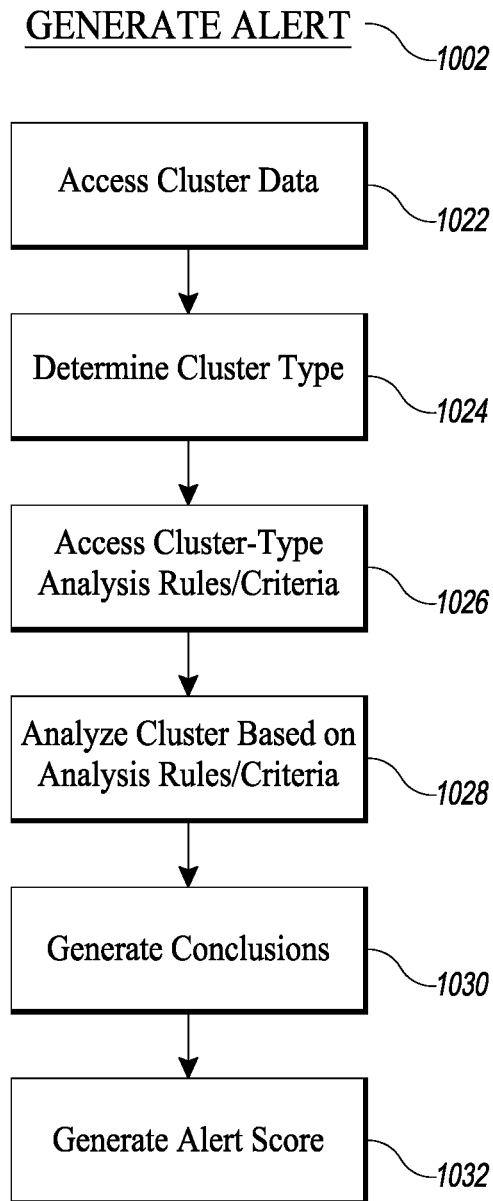


FIG. 10B

Internal Phishing	<ul style="list-style-type: none"> • <n> Senders sent emails reported to Abuse with subjects similar to "<most common subject>" • These email addresses sent <k> emails to the bank between <TIME1> and <TIME2> • There were <x> recipients; highest band was <y> • <z> domains were extracted, with <m> likely clickers identified
External Phishing	<ul style="list-style-type: none"> • <n> Senders sent emails reported to Abuse with subjects similar to "<most common subject>" • This campaign consists of <m> emails submitted to external Abuse
Internal Threat Intel	<ul style="list-style-type: none"> • '<Malware.exe>' was uploaded and calls out to <n> URLs • <x> hosts made connections to those exact URLs, with <y> more making connections to those domains/IPs • <z>% of the proxy traffic was blocked, and the last connection made was on <DATE1> • Proxy categorized <a>% of the traffic as '<Malicious/Botnet>'
External Threat Intel	<ul style="list-style-type: none"> • Domain <x> was blacklisted by <BLACKLIST1> and <y> more lists • <n> employees made <m> connections to this domain between <HOUR1> and <HOUR2>; Highest band was <z> • <k>% of proxy traffic was blocked, and <l>% was marked as malicious by Proxy
IDS (Intrusion Detection System)	<ul style="list-style-type: none"> • <n> Outbound/Inbound IDS Reports To/From <IPADDR> • Registered to <ORG> based in <COUNTRY> • Triggered <n> reports across <m> hosts • Reports span <k> hours, <j> minutes starting at <TIME1> • Most common signature was <SIGNATURE>

FIG. 10C


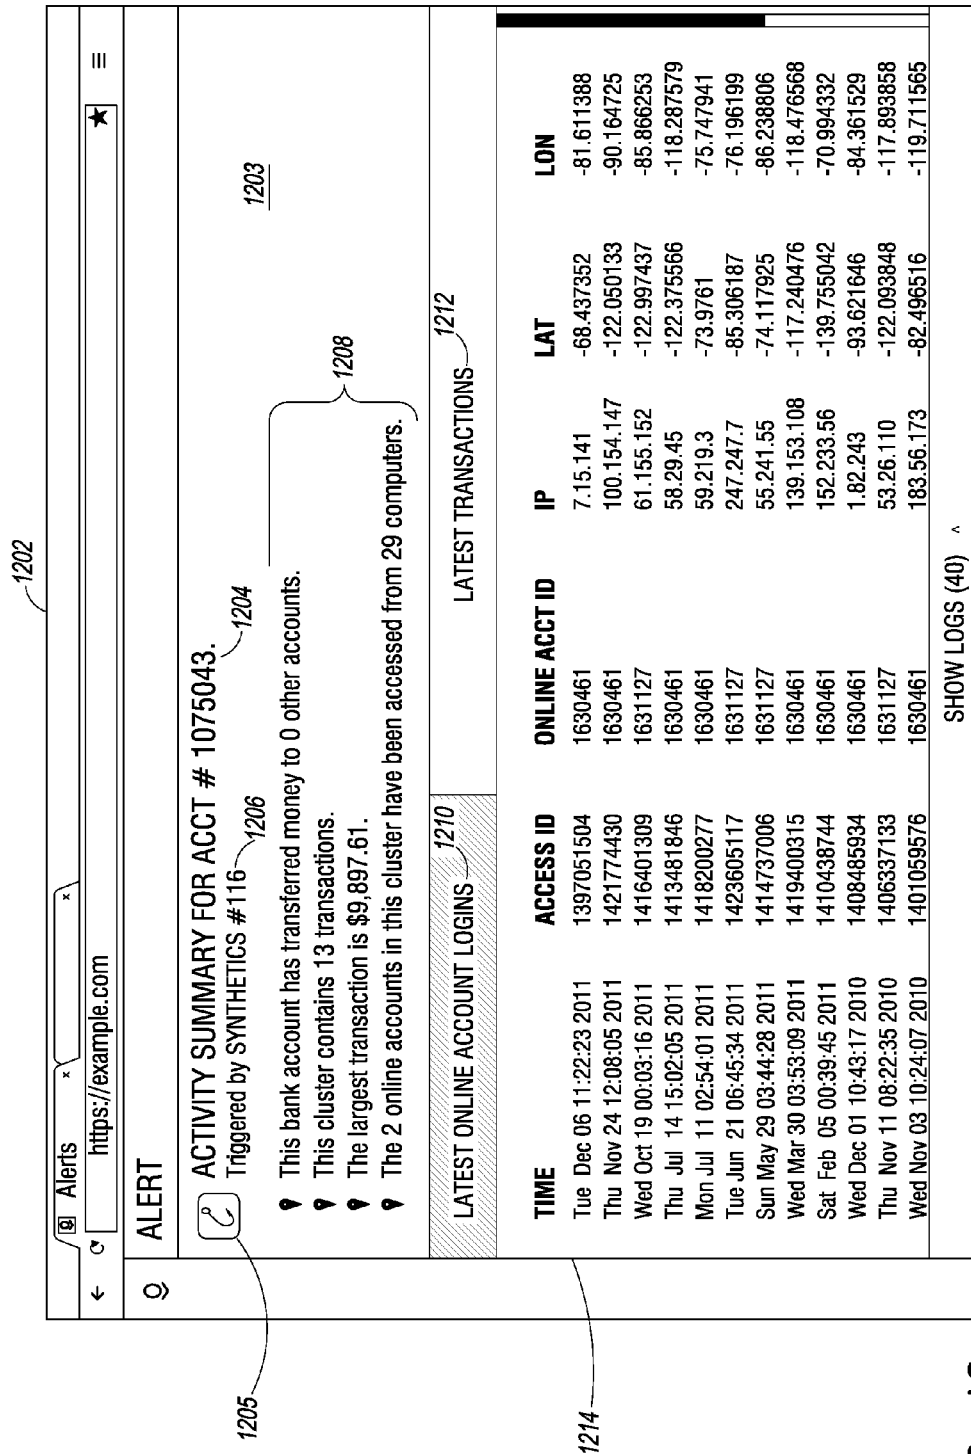
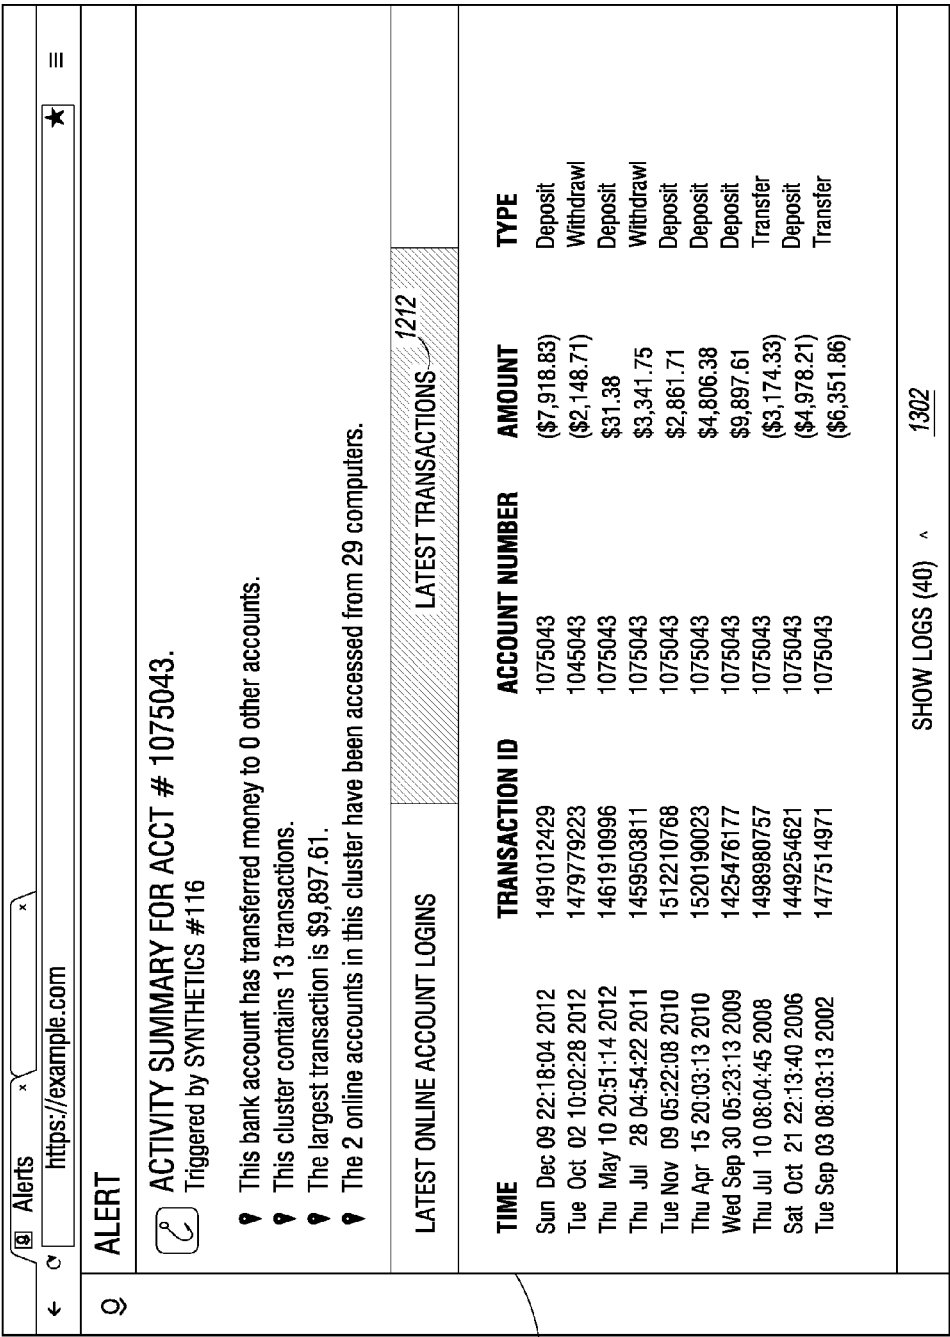


FIG. II





1214

FIG. 13

1402

Alerts x <https://example.com> ★

ALERT

ACTIVITY SUMMARY FOR ACCT # 1075043.
Triggered by SYNTHETICS #116

1406 ALL SOURCES 1408 HIDE LOGS (49) v 1302 2014 APR 30, 21:22:49

ADDRESS	CUSTOMER ID	STREET ADDRESS	CITY	STATE	ZIP	PHON		
1969 DEC 31, 16:00:00	1896665193	4921 Corbin Branch Road	Chattanooga	TN	37421	926-		
1969 DEC 31, 16:00:00	1896665193	1724 Elk Avenue	Lansing	OH	43324	430-		
TRANSACTION	TRANSACTION ID	ACCOUNT NUMBER	AMOUNT	TYPE				
2001 FEB 15, 15:05:23	1430216189	1075043	4187.46	Deposit				
2002 MAR 30, 12:36:33	1463652403	1075043	-4502.98	Withdrawal				
2002 SEP 03, 05:32:36	1452292037	1075043	-6729.79	Transfer				
2002 SEP 03, 08:03:13	1477514931	1075043	-63513.86	Transfer				
ACCOUNT	ACCOUNT NUMBER	CUSTOMER ID	ACCOUNT TYPE					
2003 MAR 05, 00:00:00	1075043	1896665193	PeRoth IRA					
ONLINE ACCOUNT	CUSTOMER ID	ONLINE ACCOUNT ID	USERNAME	EMAIL				
2003 AUG 12, 00:00:00	1896665193	1630461731	h4f6VKVbn33sofe	h4f6VKVbn33sofe@hotmailpepper.jp				
CUSTOMER	CUSTOMER ID	FIRST NAME	LAST NAME	SSN				
2003 NOV 30, 00:00:00	1896665193	GITA	SALTERS	582670007				

49 OF 49 LOGS SHOWN

1404

FIG. 14



FIG. 15

Alerts

https://example.com

★

ALERT

ACTIVITY SUMMARY FOR ACCT # 1075043.
Triggered by SYNTHETICS #116

TRANSACTION ▾

2001 FEB 15, 15:03:23

HIDE LOGS (13) ▾

2018 OCT 10, 11:40:41

TRANSACTION	TRANSACTION ID	ACCOUNT NUMBER	AMOUNT	TYPE
2001 FEB 15, 15:05:23	1430216189	1075043	4187.46	Deposit
2002 MAR 30, 12:36:33	1463652403	1075043	-4502.98	Withdrawal
2002 SEP 03, 05:32:56	1452292037	1075043	-6729.79	Transfer
2002 SEP 03, 08:03:13	1477514971	1075043	-63513.86	Transfer
2006 OCT 21, 22:13:40	1449254621	1075043	-4978.21	Deposit
2008 JUL 10, 08:04:45	1498980757	1075043	-3174.31	Transfer
2009 SEP 30, 05:23:13	1425476177	1075043	9897.61	Deposit
2010 APR 15, 20:03:13	1528190023	1075043	4806.38	Deposit

8 OF 13 LOGS SHOWN

FIG. 16

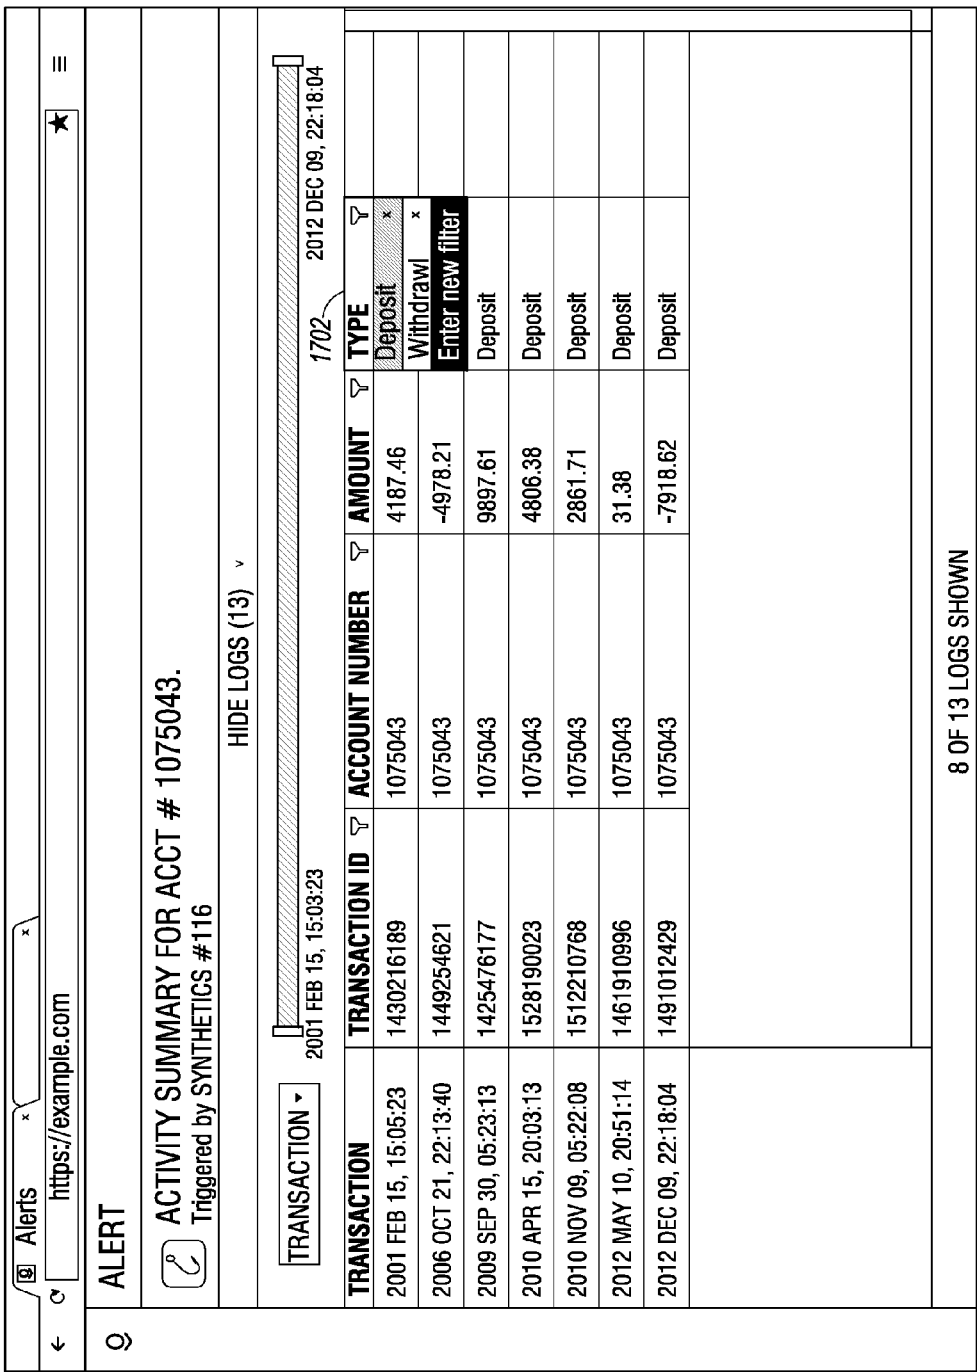


FIG. 17

1802

Alerts x
https://example.com

≡

ALERT ITEMS

All	186	← Synthetics	□ 2 Alerts Selected	Archive 1808	
Tax Fraud	20		☑ !! Activity summary for Acct # 1074911	1806	Synthetics
Beaconing Malware	15		☑ !! Activity summary for Acct # 1078704	1806	Synthetics
User-Agent Trends	12				
	5				
Internal Phishing	35		□ !! Activity summary for Acct # 1077661		Synthetics
External Phishing	40		□ !! Activity summary for Acct # 1075043		Synthetics
Internal Threats	5		□ !! Activity summary for Acct # 1075587		Synthetics
External Threats	10		□ Activity summary for Acct # 1074328		Synthetics
IDS	17		□ Activity summary for Acct # 1074320		Synthetics
Synthetics	12		□ Activity summary for Acct # 1077940		Synthetics
AML	15		□ Activity summary for Acct # 1078779		Synthetics
			□ Activity summary for Acct # 1075896		Synthetics

FIG. 18




FIG. 19




FIG. 20

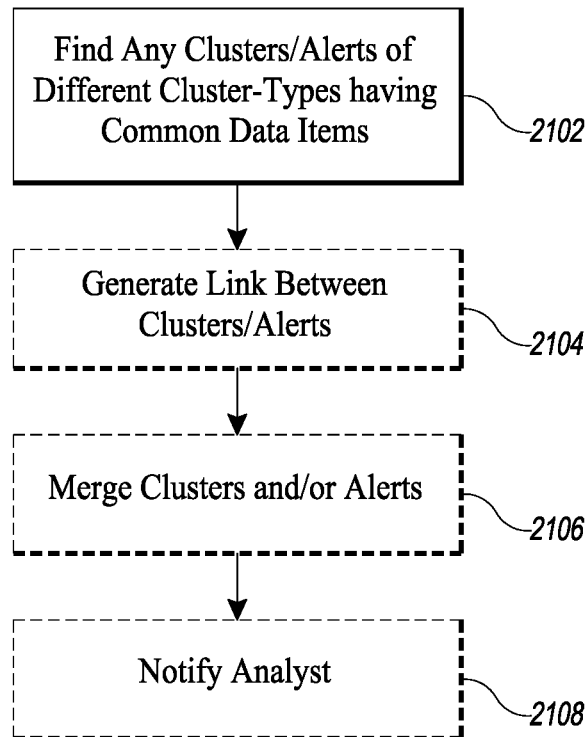



FIG. 21



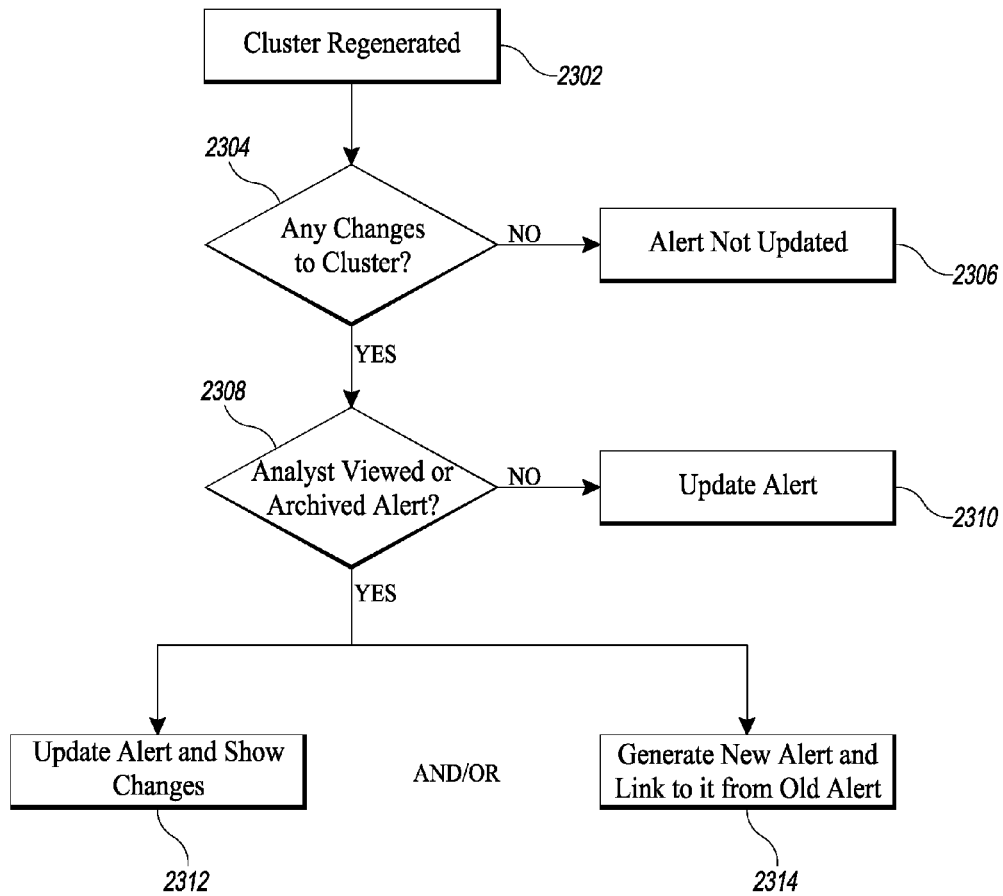


FIG. 23

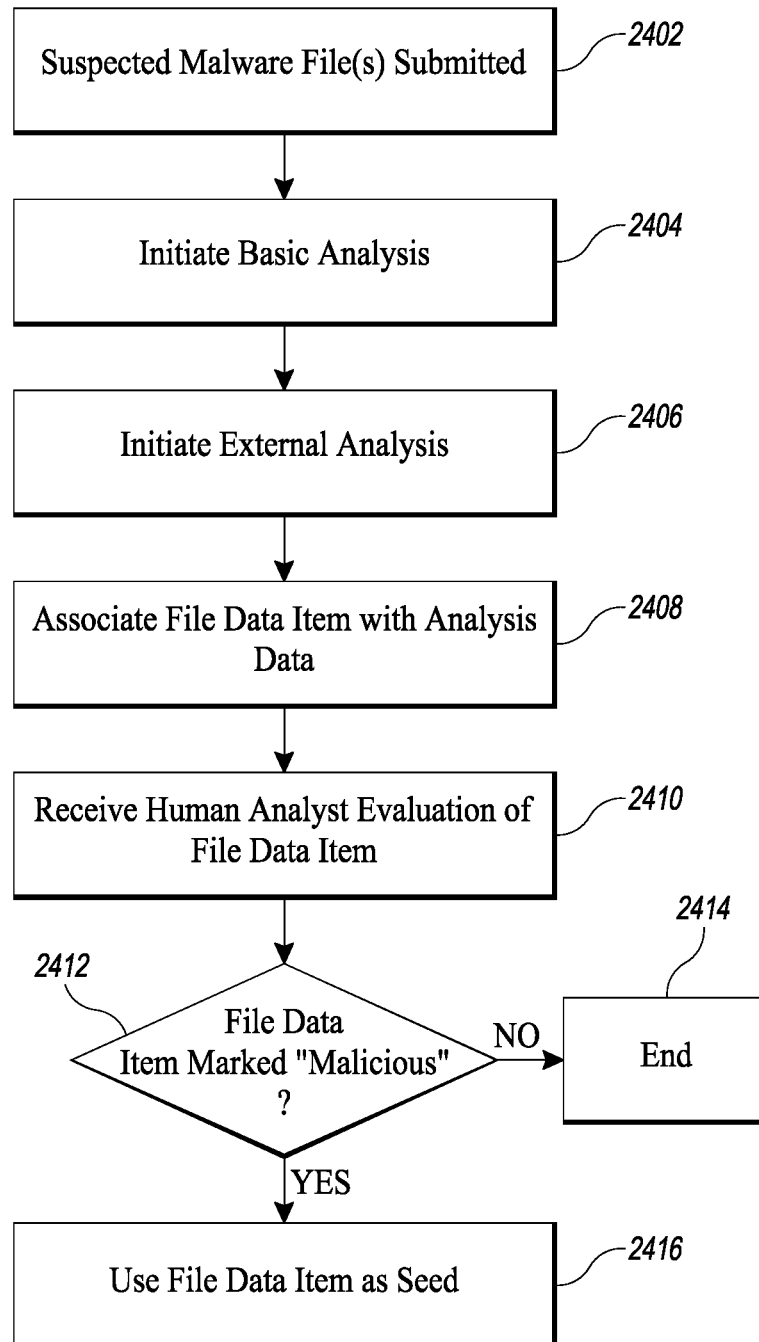


FIG. 24A

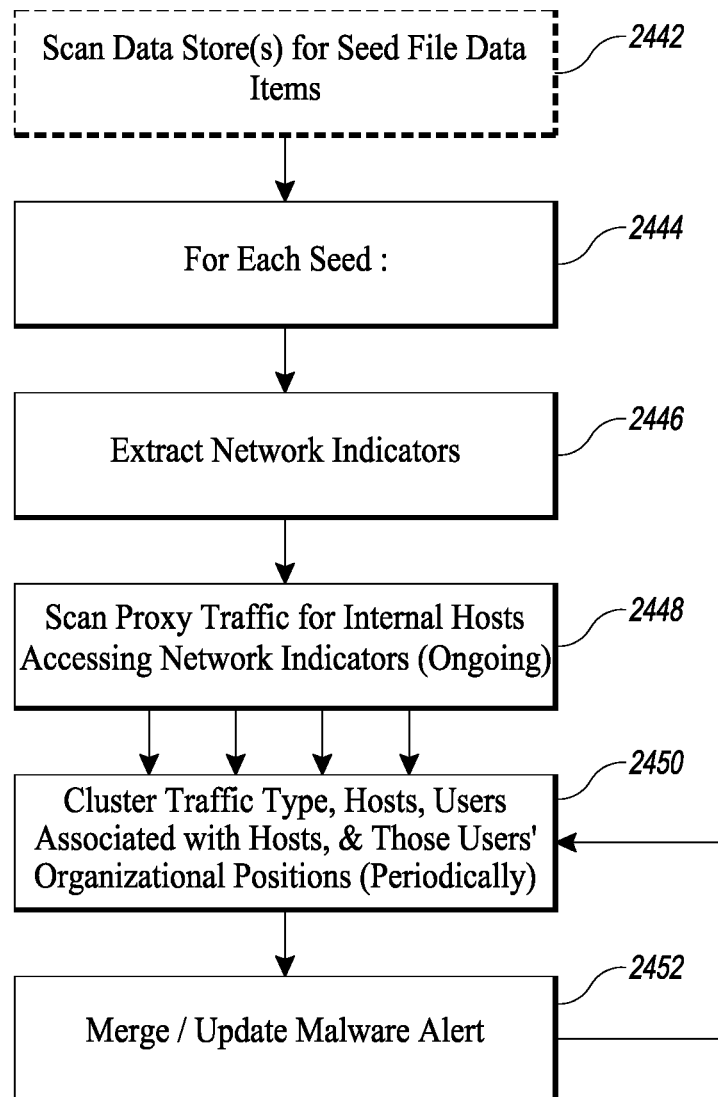


FIG. 24B

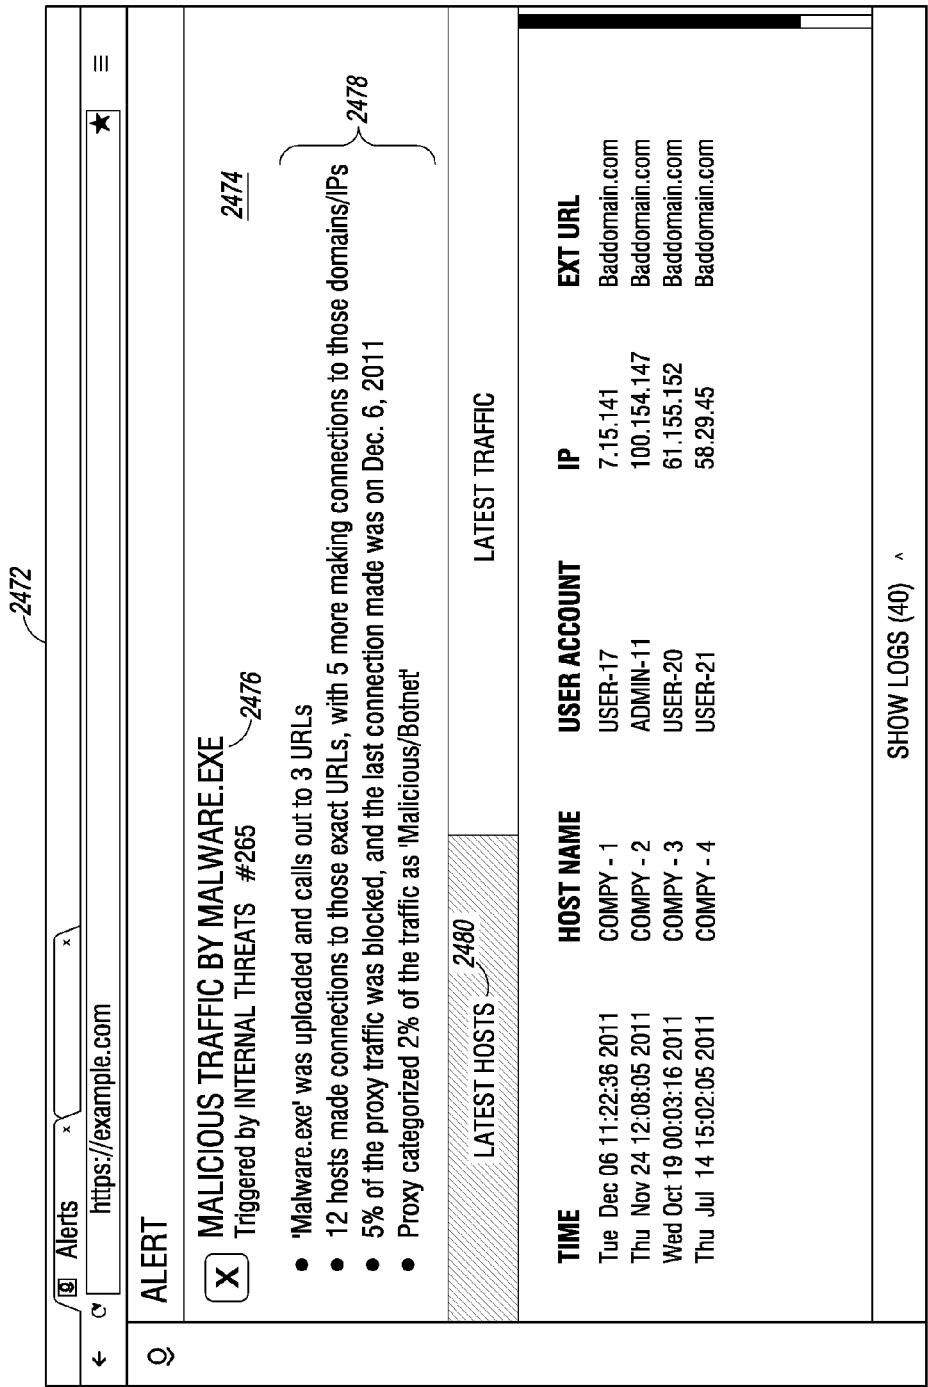


FIG. 24C

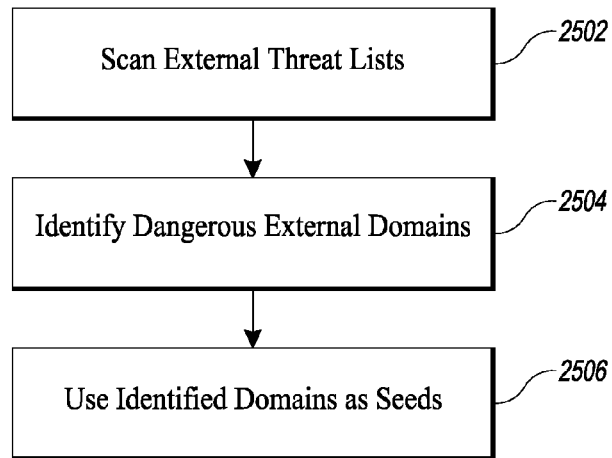


FIG. 25A

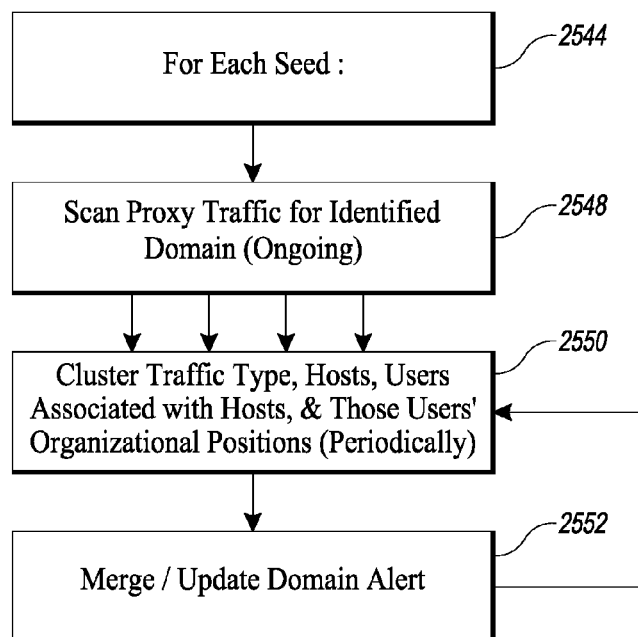


FIG. 25B

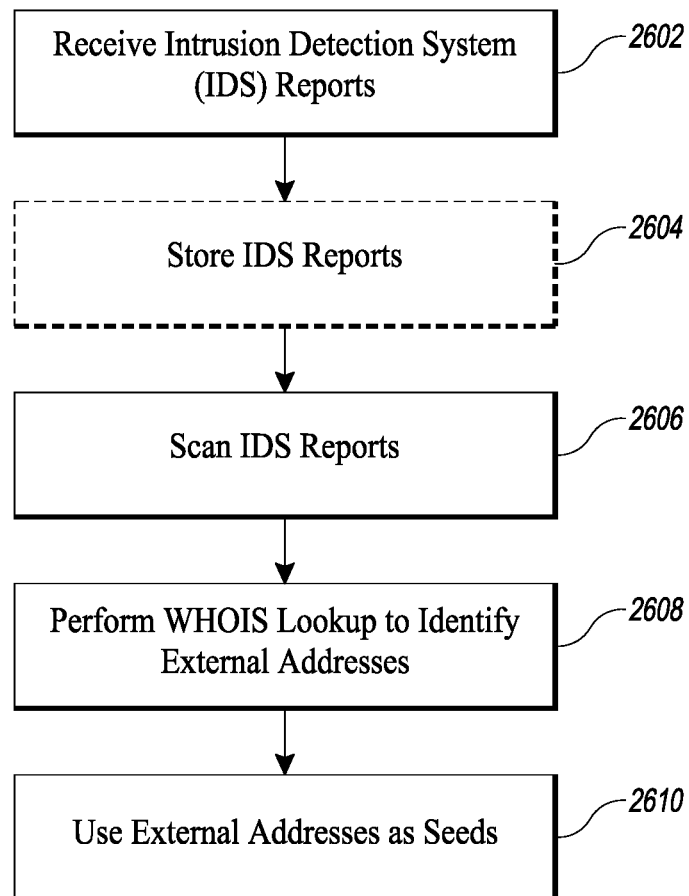


FIG. 26A

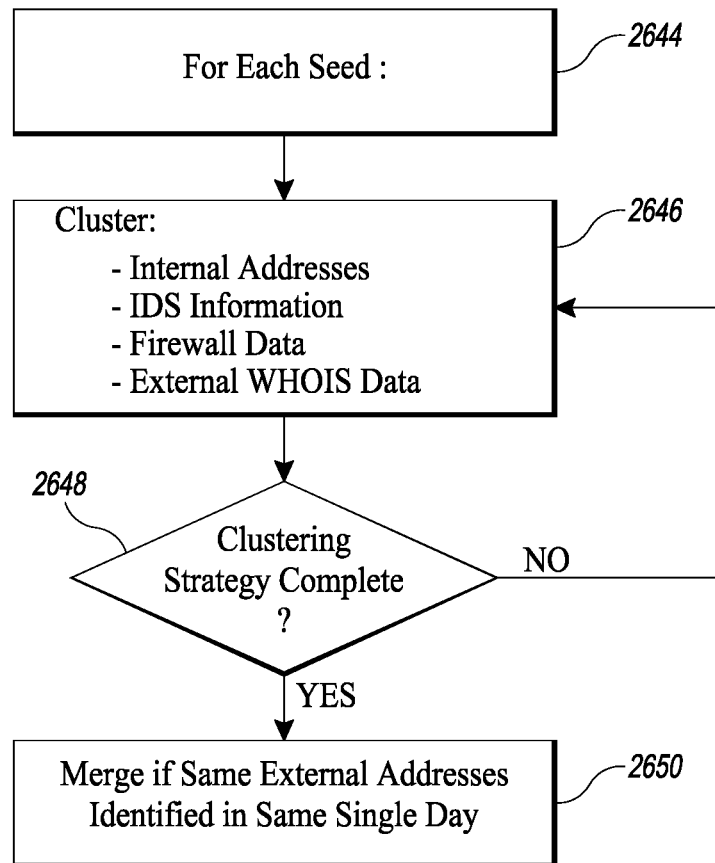


FIG. 26B




FIG. 26C

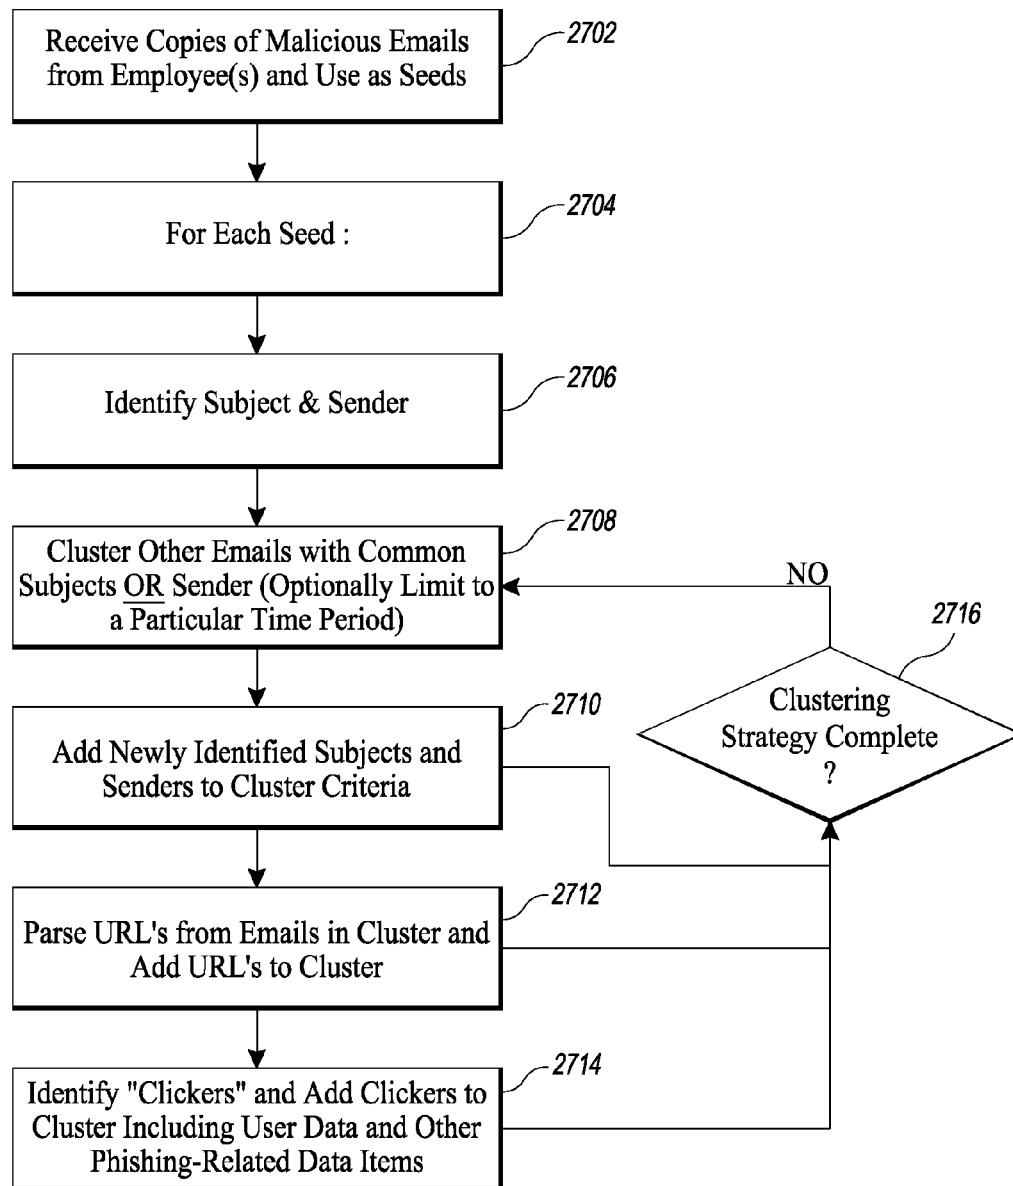


FIG. 27A

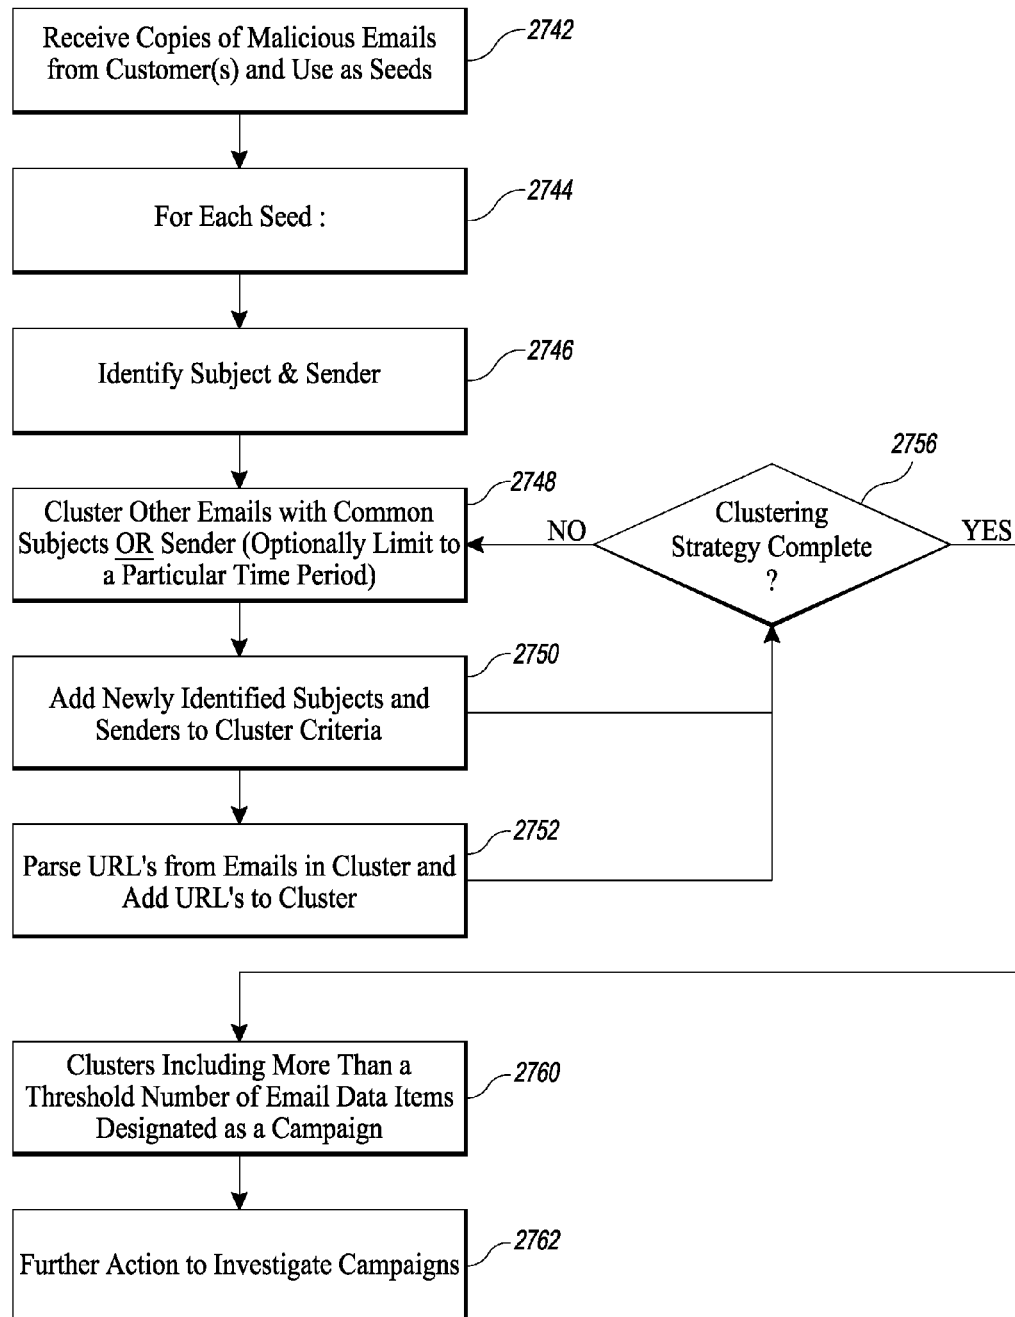


FIG. 27B

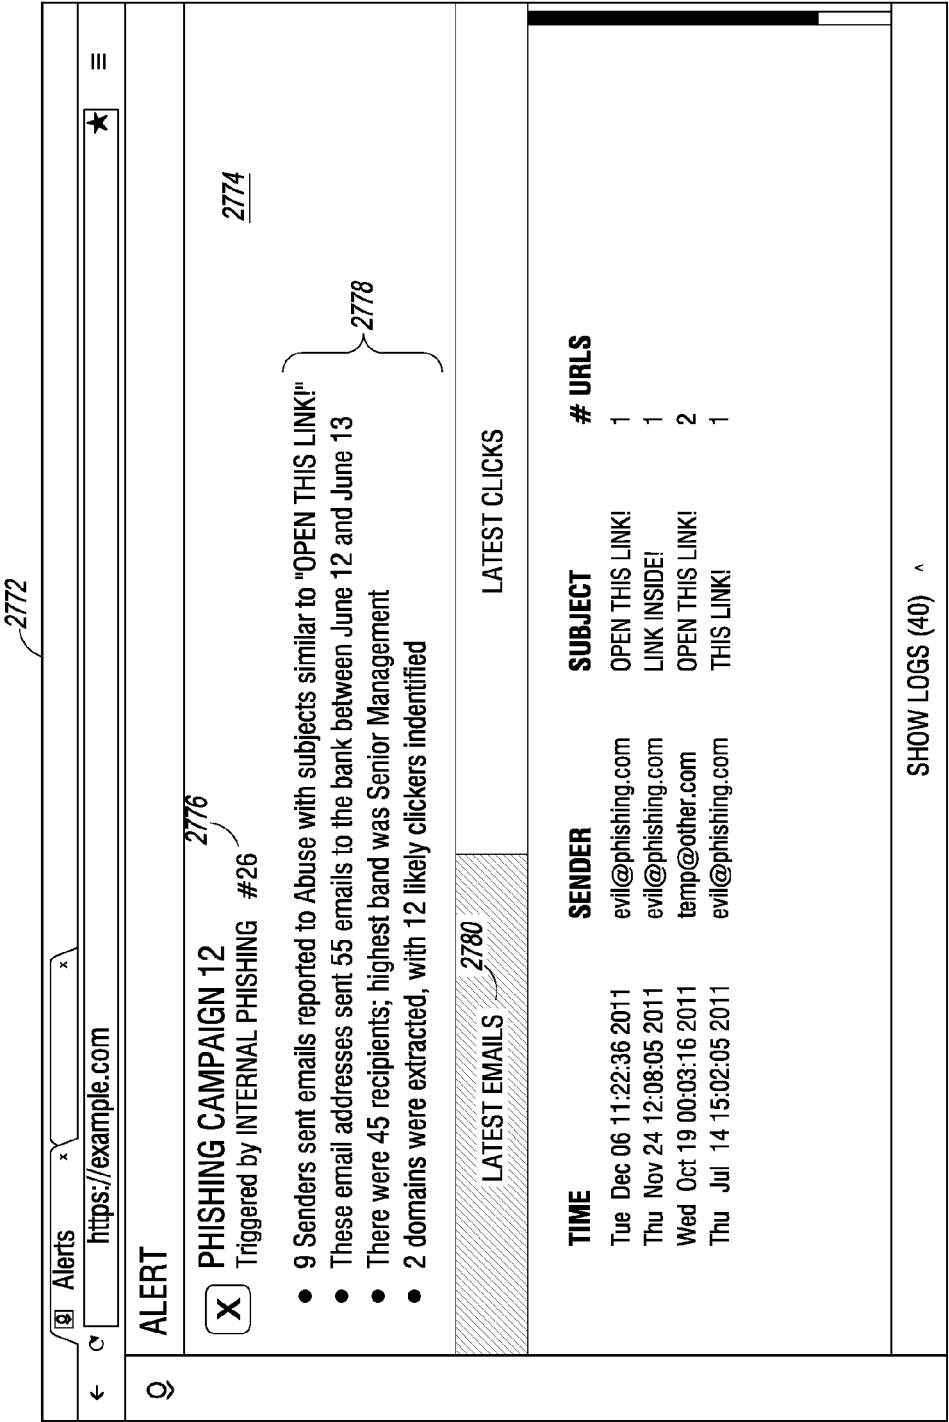


FIG. 27C

INTERNAL MALWARE DATA ITEM CLUSTERING AND ANALYSIS

CROSS-REFERENCE TO RELATED APPLICATIONS

Any and all applications for which a foreign or domestic priority claim is identified in the Application Data Sheet as filed with the present application are hereby incorporated by reference under 37 CFR 1.57.

This application is a continuation of U.S. patent application Ser. No. 14/473,920, filed Aug. 29, 2014, titled "EXTERNAL MALWARE DATA ITEM CLUSTERING AND ANALYSIS," which application claims benefit of U.S. Provisional Patent Application No. 62/020,876, filed Jul. 3, 2014, titled "DATA ITEM CLUSTERING AND ANALYSIS." The entire disclosure of each of the above items is hereby made part of this specification as if set forth fully herein and incorporated by reference for all purposes, for all that it contains.

This application is also related to the following U.S. patent applications:

Docket No.	Serial No.	Title	Filed
PALAN.235A1	13/968265	GENERATING DATA CLUSTERS WITH CUSTOMIZABLE ANALYSIS STRATEGIES	Mar. 15, 2013
PALAN.235A2	13/968213	PRIORITIZING DATA CLUSTERS WITH CUSTOMIZABLE SCORING STRATEGIES	Mar. 15, 2013
PALAN.235A1P1	14/139628	TAX DATA CLUSTERING	Dec. 23, 2013
PALAN.235A1P2	14/139603	MALWARE DATA CLUSTERING	Dec. 23, 2013
PALAN.235A1P3	14/139713	USER-AGENT DATA CLUSTERING	Dec. 23, 2013
PALAN.235A1P4	14/139640	TREND DATA CLUSTERING	Dec. 23, 2013
PALAN.245A	14/251485	FRAUD DETECTION AND SCORING	Apr. 11, 2014
PALAN.249P1	14/278963	CLUSTERING DATA BASED ON INDICATIONS OF FINANCIAL MALFEASANCE	May 15, 2015
PALAN.258PR	62/036519	MONEY LAUNDERING DETECTION AND SCORING	Aug. 12, 2014
PALAN.259A	14/473860	MALWARE DATA ITEM ANALYSIS	Aug. 29, 2014

The entire disclosure of each of the above items is hereby made part of this specification as if set forth fully herein and incorporated by reference for all purposes, for all that it contains.

BACKGROUND

Embodiments of the present disclosure generally related to automatic generation of memory-efficient clustered data structures and automatic analysis of those generated clustered data structures.

In a fraud investigation an analyst may have to make decisions regarding selection of electronic data items within an electronic collection of data. Such a collection of data may include a large number of data items that may or may not be related to one another, and which may be stored in an elec-

tronic data store or memory. For example, such a collection of data may include hundreds of thousands, millions, tens of millions, hundreds of millions, or even billions of data items, and may consume significant storage and/or memory. Determination and selection of relevant data items within such a collection of data may be extremely difficult for the analyst. Further, processing of such a large collection of data (for example, as an analyst uses a computer to sift and/or search through huge numbers of data items) may be extremely inefficient and consume significant processing and/or memory resources.

SUMMARY

The systems, methods, and devices described herein each have several aspects, no single one of which is solely responsible for its desirable attributes. Without limiting the scope of this disclosure, several non-limiting features will now be discussed briefly.

Embodiments of the present disclosure relate to a data analysis system that may automatically generate memory-efficient clustered data structures, automatically analyze those clustered data structures, and provide results of the automated analysis in an optimized way to an analyst. The automated analysis of the clustered data structures (also referred to herein as "data item clusters," "data clusters," or simply "clusters") may include an automated application of various criteria or rules so as to generate a compact, human-readable analysis of the data clusters. The human-readable analyses (also referred to herein as "conclusions" or "summaries") of the data clusters may be organized into an interactive user interface so as to enable an analyst to quickly navigate among information associated with various data clusters and efficiently evaluate those data clusters in the context of, for example, a fraud investigation. Embodiments of the present disclosure also relate to automated scoring of the clustered data structures. The interactive user interface may be updated based on the scoring, directing the human analyst to more critical data clusters (for example, data clusters more likely to be associated with fraud) in response to the analyst's inputs.

As described below, data clusters may include one or more data items. A data item may include any data, information, or things, such as a person, a place, an organization, an account, a computer, an activity, and event, and/or the like. In an example application, a human analyst may be tasked with deciding whether an account data item represents a fraudulent bank account. However, an individual data item oftentimes includes insufficient information for the analyst to make such decisions. Rather, the analyst may make better decisions based upon a collection of related data items. For instance, two financial transactions may be related by an identical account identifier or two accounts belonging to one customer may be related by an identical customer identifier or other attribute (e.g., a shared phone number or address). Some currently available systems assist the analyst by identifying data items that are directly related to an initial data item. For example, the analyst could initiate an investigation with a single suspicious data item or "seed," such as a fraudulent credit card account. If the analyst examined this data item by itself, then the analyst would not observe any suspicious characteristics. However, the analyst could request a list of data items related to the seed by a shared attribute, such as a customer identifier. In doing so, the analyst could discover an additional data item, such as an additional credit card account, which relates to the original fraudulent account because of a shared customer identifier. The analyst could

then mark the additional credit card account as potentially fraudulent, based upon the relationship of the shared customer identifier.

Although these currently available systems can be helpful in discovering related data items, they typically require the analyst to manually repeat the same series of searches for many investigations. Repeating the same investigation process consumes time and resources, such that there are often-times more investigations than can be performed. Thus, analysts typically prioritize investigations based upon the characteristics of the seeds. However, there may be insignificant differences between the seeds, so the analyst may not be able to determine the correct priority for investigations. For instance, the analyst could have to choose between two potential investigations based upon separate fraudulent credit card accounts. One investigation could reveal more potentially fraudulent credit card accounts than the other, and therefore could be more important to perform. Yet, the characteristics of the two original credit card accounts could be similar, so the analyst would not be able to choose the more important investigation. Without more information, prioritizing investigations, and evaluating data items, is difficult and error prone.

In contrast with these currently available systems, and as described above, according to various embodiments the data analysis system of the present disclosure automatically creates clusters of related data items, generates human-readable conclusions of the clusters, scores those clusters, and generates an interactive user interface in which, in response to inputs from the analyst, information related to the clusters may be efficiently provided to the analyst. Accordingly, the analyst may be enabled to efficiently evaluate the clusters.

Generation of the memory-efficient clustered data structures may be accomplished by automatic selection of an initial data item of interest (also referred to herein as a “seed”), adding of the initial data item to the memory-efficient clustered data structure (or, alternatively, designating the initial data item as the clustered data structure, or an initial iteration of the clustered data structure), and determining and adding one or more related data items to the cluster. In various embodiments, a generated cluster may include far fewer data items than the collection of data described above, and the data items included in the cluster may only include those data items that are relevant to a particular investigation (for example, a fraud investigation). Accordingly, in an embodiment, processing of the generated cluster may be highly efficient as compared to the collection of data described above. This may be because, for example, a given fraud investigation by an analyst (for example, as the analyst sifts and/or searches through data items of a cluster) may only require storage in memory of a single cluster data structure. Further, a number of data items in a cluster may be several orders of magnitude smaller than in the entire electronic collection of data described above because only data items related to each other are included in the cluster.

Additionally, the automated analysis and scoring of clusters (as mentioned above) may enable highly efficient evaluation of the various data clusters by a human analyst. For example, the interactive user interface is generated so as to enable an analyst to quickly view critical data clusters (as determined by the automated scoring), and then in response to analyst inputs, view and interact with the generated information (including, for example, the human-readable conclusions) associated with the clusters. In response to user inputs the user interface may be updated to display raw data associated with each of the generated clusters if the analyst desires to dive deeper into data associated with a given cluster.

In various embodiments, seeds may be automatically selected/generated according to various seed determination strategies, and clusters of related data items may be generated based on those seeds and according to cluster generation strategies (also referred to herein as “cluster strategies”). Also, as mentioned above, the system may generate a score, multiple scores, and/or metascores for each generated cluster, and may optionally rank or prioritize the generated clusters based on the generated scores and/or metascores. High priority clusters may be of greater interest to an analyst as they may contain related data items that meet particular criteria related to the analyst’s investigation. In an embodiment, the system may enable an analyst to advantageously start an investigation with a prioritized cluster including many related data items rather than a single randomly selected data item. Further, as described above, the cluster prioritization and summary generation may enable the processing requirements of the analyst’s investigation to be highly efficient as compared to processing of the huge collection of data described above. As mentioned above, this is because, for example, a given investigation by an analyst may only require storage in memory of a limited number of data items associated with a small number of clusters, and further, a number of data items in a cluster may be several orders of magnitude smaller than in the entire electronic collection of data described above because only data items related to each other are included in the cluster. Further, an analyst may not need to view many (or, alternatively, any) data items associated with a cluster to evaluate the cluster, but rather may evaluate the cluster based on the automatically generated cluster information.

According to an embodiment, a computer system comprises: one or more computer readable storage devices configured to store: a plurality of computer executable instructions; a plurality of data cluster types, each data cluster type associated with a data clustering strategy and a plurality of data cluster analysis rules; and a plurality of data clusters, each data cluster associated with a data cluster type and previously generated according to the associated respective data clustering strategy, each data cluster further including one or more data items and associated metadata; and one or more hardware computer processors in communication with the one or more computer readable storage devices and configured to execute the plurality of computer executable instructions in order to cause the computer system to, for each particular data cluster of the plurality of data clusters: access the particular data cluster from the one or more computer readable storage devices; determine the data cluster type associated with the particular data cluster; analyze the particular data cluster based on the data cluster analysis rules associated with the determined data cluster type; determine an alert score for the particular data cluster based on the analysis of the particular data cluster; and generate one or more human-readable conclusions regarding the particular data cluster.

According to another embodiment, the alert score indicates a degree of correlation between characteristics of the particular data cluster and the accessed data cluster analysis rules.

According to yet another embodiment, the degree of correlation is based on both an assessment of risk associated with the particular data cluster and a confidence level in accuracy of the assessment of risk.

According to another embodiment, a relatively higher alert score indicates a data cluster that is relatively more important for a human analyst to evaluate, and a relatively lower alert score indicates a data cluster that is relatively less important for the human analyst to evaluate.

5

According to yet another embodiment, each alert score for respective data clusters is assigned to a category indicating a high degree of correlation, a medium degree of correlation, or a low degree of correlation.

According to another embodiment, the high degree of correlation is associated with a first color, the medium degree of correlation is associated with a second color, and the low degree of correlation is associated with a third color.

According to yet another embodiment, the one or more hardware computer processors are further configured to execute the plurality of computer executable instructions in order to cause the computer system to: for each particular data cluster of the plurality of data clusters: generate an alert, the alert comprising the alert score, the one or more human-readable conclusions, the one or more data items associated with the cluster, and the metadata associated with each of the one or more data items.

According to another embodiment, the one or more hardware computer processors are further configured to execute the plurality of computer executable instructions in order to cause the computer system to: generate a user interface including a list of user-selectable alert indicators, an alert indicator being provided for each of the generated alerts, each of the alert indicators providing a summary of information associated with respective generated alerts.

According to yet another embodiment, the one or more hardware computer processors are further configured to execute the plurality of computer executable instructions in order to cause the computer system to: in response to a selection of an alert indicator by a human analyst: generate an alert display, the alert display including at least an indication of the alert score and a list of the one or more human-readable conclusions.

According to another embodiment, the alert display further includes a table of information associated with the one or more data items and associated metadata of the particular data cluster, and wherein the table of information includes a mixture of information of various types.

According to yet another embodiment, the table of information includes one or more user interface controls selectable by a human analysis in order to filter according information type and/or time period.

According to another embodiment, the one or more hardware computer processors are further configured to execute the plurality of computer executable instructions in order to cause the computing system to: for each particular data cluster of the plurality of data clusters: determine whether any other data cluster having a data cluster type different from the data cluster type of the particular data cluster is associated with a data item that is also associated with the particular data cluster.

According to yet another embodiment, the alert display further indicates and provides a link to any of the other data clusters having the data cluster type different from the data cluster type of the particular data cluster and determined to be associated with the data item that is also associated with the particular data cluster of the alert display.

According to another embodiment, a notification is provided to a human analyst via the alert display when the data cluster associated with the alert display has been regenerated such that the data cluster is changed.

According to an embodiment, a computer system comprises: one or more computer readable storage devices configured to store: a plurality of computer executable instructions; a data clustering strategy; and a plurality of data items including at least: file data items, each file data item associated with at least one suspected malware file and one or more

6

network indicators; and network-related data items associated with captured communications between an internal network and an external network, the network-related data items including at least one of: external Internet Protocol addresses, external domains, external computerized devices, internal Internet Protocol addresses, internal computerized devices, users of particular computerized devices, or organizational positions associated with users of particular computerized devices; and one or more hardware computer processors in communication with the one or more computer readable storage devices and configured to execute the plurality of computer executable instructions in order to cause the computer system to: access, from the one or more computer readable storage devices, the file data items; determine, for each file data item of the file data items, whether or not the file data item is been marked as a seed; and for each of the file data items marked as a seed, generate a data item cluster based on the data clustering strategy by at least: adding the seed to the data item cluster; identifying one or more of the network indicators that are associated with the seed; identifying one or more of the network-related data items associated with at least one of the identified one or more of the network indicators; adding, to the data item cluster, the one or more identified network-related data items; identifying an additional one or more data items, including file data items and/or network-related data items, associated with any data items of the data item cluster; and adding, to the data item cluster, the additional one or more data items.

According to another embodiment, each of the data items of the data item cluster identify at least an internal computerized device, a user of the internal computerized device, and an organizational position associated with the user.

According to yet another embodiment, the one or more hardware computer processors are further configured to execute the plurality of computer executable instructions in order to cause the one or more hardware computer processors to: scan communications between the internal network and the external network so as to identify additional network-related data items; and store the additional network-related data items in the one or more computer readable storage devices.

According to another embodiment, the communications are continuously scanned via a proxy.

According to yet another embodiment, the one or more network indicators include at least an external Internet Protocol address or an external domain.

According to another embodiment, the one or more of the network indicators that are associated with the seed comprise network indicators that are contacted by the at least one suspected malware file associated with the seed when the at least one suspected malware file is executed.

According to yet another embodiment, determining whether or not the file data item is marked as a seed comprises determining whether or not the file data item has been marked by a human analyst as a seed.

According to another embodiment, each of the file data items is processed by the computer system by at least: initiating an analysis of the file data item including the at least one suspected malware file, wherein the analysis of the file data item generates a plurality of analysis information items including at least one of calculated hashes, file properties, academic analysis information, file execution information, or third-party analysis information; associating the plurality of analysis information items with the file data item; and generating a user interface including one or more user selectable portions presenting various of the analysis information items,

the user interface usable by the human analyst to determine one or more characteristics of the file data item and to mark the file data item as a seed.

According to yet another embodiment, the file data item is marked by a human analyst as a seed via a user interface of the computer system.

According to another embodiment, initiating analysis of the file data item and generating the plurality of analysis information items comprises: initiating an internal analysis of the file data item; and initiating an external analysis of the file data item, wherein the internal analysis includes analysis performed by the one or more hardware computer processors, and wherein the internal analysis includes at least one of calculation of an MD5 hash of the file data item, calculation of a SHA-1 hash of the file data item, or calculation of a size of the file data item, and wherein the external analysis includes analysis performed by at least a second computer system, and wherein the external analysis includes execution of the file data item in a sandboxed environment and analysis of the file data item by a third-party malware analysis service.

According to yet another embodiment, the one or more computer readable storage devices are further configured to store: a plurality of data cluster analysis rules associated with the data clustering strategy, and the one or more hardware computer processors are further configured to execute the plurality of computer executable instructions in order to cause the one or more hardware computer processors to: for each generated data item cluster: access the plurality of data cluster analysis rules associated with the data clustering strategy; analyze the data item cluster based on the accessed data cluster analysis rules; and based on the analysis of the data item cluster: determine an alert score for the data item cluster; and generate one or more human-readable conclusions regarding the data item cluster.

According to another embodiment, the alert score indicates a degree of correlation between characteristics of the data item cluster and the accessed data cluster analysis rules.

According to yet another embodiment, a relatively higher alert score indicates a data cluster that is relatively more important for a human analyst to evaluate, and a relatively lower alert score indicated a data cluster that is relatively less important for the human analyst to evaluate.

According to another embodiment, each alert score for respective data clusters is assigned to a category indicating a high degree of correlation, a medium degree of correlation, or a low degree of correlation.

According to yet another embodiment, the high degree of correlation is associated with a first color, the medium degree of correlation is associated with a second color, and the low degree of correlation is associated with a third color.

According to another embodiment, the one or more hardware computer processors are further configured to execute the plurality of computer executable instructions in order to cause the computer system to: for each generated data item cluster: generate an alert, the alert comprising a the alert score, the one or more human-readable conclusions, the data items associated with the data item cluster, and metadata associated with the data items of the data item cluster.

According to yet another embodiment, the one or more hardware computer processors are further configured to execute the plurality of computer executable instructions in order to cause the computer system to: generate a user interface including a list of user-selectable alert indicators, an alert indicator being provided for each of the generated alerts, each of the alert indicators providing a summary of information associated with respective generated alerts.

According to another embodiment, the one or more hardware computer processors are further configured to execute the plurality of computer executable instructions in order to cause the computer system to: in response to a selection of an alert indicator by a human analyst: generate an alert display, the alert display including at least an indication of the alert score and a list of the one or more human-readable conclusions.

According to yet another embodiment, the one or more human-readable conclusions each comprise a phrase or sentence including one or more indications of summary or aggregated data associated with a plurality of the data items of the data item cluster.

According to another embodiment, generating the one or more human-readable conclusions comprises: selecting, based on the data cluster type associated with the particular data cluster, one or more conclusion templates; and populating the one or more conclusion templates with data associated with the particular data cluster.

According to an embodiment, a computer system comprises: one or more computer readable storage devices configured to store: a plurality of computer executable instructions; a data clustering strategy; and a plurality of data items including at least: external domain data items; and network-related data items associated with captured communications between an internal network and an external network, the network-related data items including at least one of: external Internet Protocol addresses, external domains, external computerized devices, internal Internet Protocol addresses, internal computerized devices, users of particular computerized devices, or organizational positions associated with users of particular computerized devices; and one or more hardware computer processors in communication with the one or more computer readable storage devices and configured to execute the plurality of computer executable instructions in order to cause the computer system to: scan one or more threat lists stored external to the internal network, each of the threat lists including information related to previously identified malware threats and information related to those previously identified malware threats including external domain data items; identify one or more external domain data items included in the one or more threat lists, each of the one or more external domain data items being associated with a malicious domain; designate each of the identified one or more external domain data items as a seed; and for each of the designated seeds, generate a data item cluster based on the data clustering strategy by at least: adding the seed to the data item cluster; identifying one or more of the network-related data items associated with the seed; adding, to the data item cluster, the one or more identified network-related data items; identifying an additional one or more data items, including external domain data items and/or network-related data items, associated with any data items of the data item cluster; and adding, to the data item cluster, the additional one or more data items.

According to another embodiment, data items of the data item cluster include at least an internal computerized device, a user of the internal computerized device, and an organizational position associated with the user.

According to yet another embodiment, the one or more hardware computer processors are further configured to execute the plurality of computer executable instructions in order to cause the one or more hardware computer processors to: scan communications between the internal network and the external network so as to generate additional network-related data items; and store the additional network-related data items in the one or more computer readable storage devices.

According to another embodiment, the communications are continuously scanned via a proxy.

According to yet another embodiment, identifying the one or more of the network-related data items associated with the seed comprises: for each of the network-related data items, determining whether the network-related data item is associated with a communication with the malicious domain associated with the seed; and in response to determining that the network-related data item is associated with a communication with the malicious domain associated with the seed, identifying the network-related data item as being associated with the seed.

According to another embodiment, the data item clusters are continuously updated.

According to yet another embodiment, the one or more computer readable storage devices are further configured to store: a plurality of data cluster analysis rules associated with the data clustering strategy, and the one or more hardware computer processors are further configured to execute the plurality of computer executable instructions in order to cause the one or more hardware computer processors to: for each generated data item cluster: access the plurality of data cluster analysis rules associated with the data clustering strategy; analyze the data item cluster based on the accessed data cluster analysis rules; and based on the analysis of the data item cluster: determine an alert score for the data item cluster; and generate one or more human-readable conclusions regarding the data item cluster.

According to another embodiment, the alert score indicates a degree of correlation between characteristics of the data item cluster and the accessed data cluster analysis rules.

According to yet another embodiment, the degree of correlation is based on both an assessment of risk associated with the particular data cluster and a confidence level in accuracy of the assessment of risk.

According to another embodiment, a relatively higher alert score indicates a data cluster that is relatively more important for a human analyst to evaluate, and a relatively lower alert score indicated a data cluster that is relatively less important for the human analyst to evaluate.

According to yet another embodiment, each alert score for respective data clusters is assigned to a category indicating a high degree of correlation, a medium degree of correlation, or a low degree of correlation.

According to another embodiment, the high degree of correlation is associated with a first color, the medium degree of correlation is associated with a second color, and the low degree of correlation is associated with a third color.

According to yet another embodiment, the one or more hardware computer processors are further configured to execute the plurality of computer executable instructions in order to cause the computer system to: for each generated data item cluster: generate an alert, the alert comprising a the alert score, the one or more human-readable conclusions, the data items associated with the data item cluster, and metadata associated with the data items of the data item cluster.

According to another embodiment, the one or more hardware computer processors are further configured to execute the plurality of computer executable instructions in order to cause the computer system to: generate a user interface including a list of user-selectable alert indicators, an alert indicator being provided for each of the generated alerts, each of the alert indicators providing a summary of information associated with respective generated alerts.

According to yet another embodiment, the one or more hardware computer processors are further configured to execute the plurality of computer executable instructions in

order to cause the computer system to: in response to a selection of an alert indicator by a human analyst: generate an alert display, the alert display including at least an indication of the alert score and a list of the one or more human-readable conclusions.

According to another embodiment, the alert display further includes a table of information associated with the data items associated with the data item cluster of the alert, and metadata associated with the data items of the data item cluster of the alert.

According to yet another embodiment, the table of information includes a mixture of information of various types.

According to another embodiment, the table of information includes one or more user interface controls selectable by a human analysis in order to filter according information type and/or time period.

According to yet another embodiment, the one or more human-readable conclusions each comprise a phrase or sentence including one or more indications of summary or aggregated data associated with a plurality of the data items of the data item cluster.

According to another embodiment, generating the one or more human-readable conclusions comprises: selecting, based on the data cluster type associated with the particular data cluster, one or more conclusion templates; and populating the one or more conclusion templates with data associated with the particular data cluster.

According to an embodiment, a computer system comprises: one or more computer readable storage devices configured to store: a plurality of computer executable instructions; a data clustering strategy; and a plurality of data items including at least: intrusion detection system reports, each intrusion detection system report associated with at least a source Internet Protocol address and a destination Internet Protocol address; and network-related data items associated with captured communications between an internal network and an external network, the network-related data items including at least one of: external Internet Protocol addresses, external domains, external computerized devices, internal Internet Protocol addresses, internal computerized devices, users of particular computerized devices, intrusion detection system information, network firewall data, or WHOIS information; and one or more hardware computer processors in communication with the one or more computer readable storage devices and configured to execute the plurality of computer executable instructions in order to cause the computer system to: receive an intrusion detection system report including a communication between a source Internet Protocol address and a destination Internet Protocol address; initiate an automated lookup to determine which of the source Internet Protocol address and the destination Internet Protocol address is an external Internet Protocol address, the external Internet Protocol address being external to the internal network; designate the external Internet Protocol address as a seed; and generate a data item cluster based on the data clustering strategy by at least: adding the seed to the data item cluster; identifying one or more of the network-related data items associated with the seed; and adding, to the data item cluster, the one or more identified network-related data items.

According to another embodiment, generating the data item cluster based on the data clustering strategy further comprises: identifying additional one or more data items associated with any data items of the data item cluster; and adding, to the data item cluster, the additional one or more data items.

According to yet another embodiment, the data item cluster includes at least the source Internet Protocol address, the

11

destination Internet Protocol address, an internal computerized device associated with an Internet Protocol address in the cluster, a user of the internal computerized device, and WHOIS information associated with the external Internet Protocol address.

According to another embodiment, the one or more hardware computer processors are further configured to execute the plurality of computer executable instructions in order to cause the one or more hardware computer processors to: receive a second intrusion detection system report including a communication between a second source Internet Protocol address and a second destination Internet Protocol address; initiate an automated lookup to determine which of the second source Internet Protocol address and the second destination Internet Protocol address is a second external Internet Protocol address being external to the internal network; compare the external Internet Protocol address to the second external Internet Protocol address; in response to determining, based on the comparison, that the external Internet Protocol address and the second external Internet Protocol address are the same, add the second external Internet Protocol address to the data item cluster; and in response to determining, based on the comparison, that the external Internet Protocol address and the second external Internet Protocol address are not the same: designate the second external Internet Protocol address as a second seed; and generate a second data item cluster based on the data clustering strategy and the second seed.

According to yet another embodiment, the second external Internet Protocol address is added to the data item cluster only if the intrusion detection system report and the second intrusion detection system report are received on a same day.

According to another embodiment, the one or more hardware computer processors are further configured to execute the plurality of computer executable instructions in order to cause the one or more hardware computer processors to: scan communications between the internal network and the external network so as to generate additional network-related data items; and store the additional network-related data items in the one or more computer readable storage devices.

According to yet another embodiment, received intrusion detection system reports are automatically stored in the one or more computer readable storage devices, and the one or more hardware computer processors are further configured to execute the plurality of computer executable instructions in order to cause the one or more hardware computer processors to: identify newly received intrusion detection system reports; initiate automated lookups to determine external Internet Protocol addresses associated with each of the newly received intrusion detection system reports; designate the determined external Internet Protocol addresses as seeds; and generate data item clusters based on the data clustering strategy and the seeds.

According to another embodiment, data item clusters generated based on common external Internet Protocol addresses are merged.

According to yet another embodiment, data item clusters generated based on common external Internet Protocol addresses are merged only if the associated intrusion detection system reports are received on a same day.

According to another embodiment, the one or more computer readable storage devices are further configured to store: a plurality of data cluster analysis rules associated with the data clustering strategy, and the one or more hardware computer processors are further configured to execute the plurality of computer executable instructions in order to cause the one or more hardware computer processors to: for the gener-

12

ated data item cluster: access the plurality of data cluster analysis rules associated with the data clustering strategy; analyze the data item cluster based on the accessed data cluster analysis rules; and based on the analysis of the data item cluster: determine an alert score for the data item cluster; and generate one or more human-readable conclusions regarding the data item cluster.

According to yet another embodiment, the alert score indicates a degree of correlation between characteristics of the data item cluster and the accessed data cluster analysis rules.

According to another embodiment, the degree of correlation is based on both an assessment of risk associated with the particular data cluster and a confidence level in accuracy of the assessment of risk.

According to yet another embodiment, a relatively higher alert score indicates a data cluster that is relatively more important for a human analyst to evaluate, and a relatively lower alert score indicated a data cluster that is relatively less important for the human analyst to evaluate.

According to another embodiment, each alert score for respective data clusters is assigned to a category indicating a high degree of correlation, a medium degree of correlation, or a low degree of correlation.

According to yet another embodiment, the high degree of correlation is associated with a first color, the medium degree of correlation is associated with a second color, and the low degree of correlation is associated with a third color.

According to another embodiment, the one or more hardware computer processors are further configured to execute the plurality of computer executable instructions in order to cause the computer system to: for the generated data item cluster: generate an alert, the alert comprising a the alert score, the one or more human-readable conclusions, the data items associated with the data item cluster, and metadata associated with the data items of the data item cluster.

According to yet another embodiment, the one or more hardware computer processors are further configured to execute the plurality of computer executable instructions in order to cause the computer system to: generate a user interface including a list of user-selectable alert indicators, an alert indicator being provided for each of the generated alerts, each of the alert indicators providing a summary of information associated with respective generated alerts.

According to another embodiment, the one or more hardware computer processors are further configured to execute the plurality of computer executable instructions in order to cause the computer system to: in response to a selection of an alert indicator by a human analyst: generate an alert display, the alert display including at least an indication of the alert score and a list of the one or more human-readable conclusions.

According to yet another embodiment, the one or more human-readable conclusions each comprise a phrase or sentence including one or more indications of summary or aggregated data associated with a plurality of the data items of the data item cluster.

According to another embodiment, generating the one or more human-readable conclusions comprises: selecting, based on the data cluster type associated with the particular data cluster, one or more conclusion templates; and populating the one or more conclusion templates with data associated with the particular data cluster.

According to an embodiment, a computer system comprises: one or more computer readable storage devices configured to store: a plurality of computer executable instructions; a data clustering strategy; and a plurality of data items including at least: email data items, each of the email data

13

items including at least a subject and a sender, each of the email data items potentially associated with phishing activity; and phishing-related data items related to a communications network of an organization, the phishing-related data items including at least one of: internal Internet Protocol addresses of the communications network, computerized devices of the communications network, users of particular computerized devices, organizational positions associated with users of particular computerized devices, or URLs and/or external domains visited by users of particular computerized devices; and one or more hardware computer processors in communication with the one or more computer readable storage devices and configured to execute the plurality of computer executable instructions in order to cause the computer system to: access an email data item transmitted to one or more of the users of respective computerized devices within the network of the organization, the email data item including at least a subject and a sender; designate the received email data item as a seed; and generate a data item cluster based on the data clustering strategy by at least: adding the seed to the data item cluster; determining the subject and the sender associated with the seed; identifying one or more of the email data items having a same subject as the determined subject or a same sender as the determined sender; adding the identified one or more email data items to the data item cluster; adding one or more parsed URLs from the email data items to the data item cluster; identifying one or more users who are both recipients of at least one of the email data items of the data item cluster and visitors of one of the URLs of the data item cluster; and adding the identified one or more users, including data related to the one or more users, to the data item cluster.

According to another embodiment, generating the data item cluster based on the data clustering strategy further comprises: determining any new subjects or new senders associated with email data items of the data item cluster that are different from the determined subjects or the determined senders; identifying a second one or more of the email data items having a same subject as the determined new subject, or a same sender as the determined new sender; and adding the identified second one or more email data items to the data item cluster.

According to yet another embodiment, the identified one or more email data items are added to the data item cluster only if received by one or more computerized devices within the network within a predetermined period of time from a time that the seed was received.

According to another embodiment, the period of time comprises at least one of a number of hours, a number of days, or a number of weeks.

According to yet another embodiment, the predetermined period of time is further determined based on other email data items in the data item cluster.

According to another embodiment, generating the data item cluster based on the data clustering strategy further comprises: identifying additional one or more data items associated with any data items of the data item cluster; and adding, to the data item cluster, the additional one or more data items.

According to yet another embodiment, identifying the one or more users further comprises: scanning communications on the communications network of the organization so as to generate phishing-related data items including URLs visited by particular users; extracting recipients of the email data items of the data item cluster associated with respective parsed URLs; and for any parsed URL matching a URL visited by a particular user, if the extracted recipient of the

14

email data item associated with the parsed URL matches the particular user, then identifying the user.

According to another embodiment, the communications are continuously scanned via a proxy.

According to yet another embodiment, the one or more hardware computer processors are further configured to execute the plurality of computer executable instructions in order to cause the one or more hardware computer processors to: continuously receive email data items from users of respective computing devices of the organization, designate the received email data items as seeds, and generate data items clusters based on the data clustering strategy.

According to another embodiment, the data related to the one or more users includes an organizational position associated with the user.

According to yet another embodiment, the one or more computer readable storage devices are further configured to store: a plurality of data cluster analysis rules associated with the data clustering strategy, and the one or more hardware computer processors are further configured to execute the plurality of computer executable instructions in order to cause the one or more hardware computer processors to: for the generated data item cluster: access the plurality of data cluster analysis rules associated with the data clustering strategy; analyze the data item cluster based on the accessed data cluster analysis rules; and based on the analysis of the data item cluster: determine an alert score for the data item cluster; and generate one or more human-readable conclusions regarding the data item cluster.

According to another embodiment, the alert score indicates a degree of correlation between characteristics of the data item cluster and the accessed data cluster analysis rules.

According to yet another embodiment, the degree of correlation is based on both an assessment of risk associated with the particular data cluster and a confidence level in accuracy of the assessment of risk.

According to another embodiment, each alert score for respective data clusters is assigned to a category indicating a high degree of correlation, a medium degree of correlation, or a low degree of correlation, and wherein the high degree of correlation is associated with a first color, the medium degree of correlation is associated with a second color, and the low degree of correlation is associated with a third color.

According to yet another embodiment, the one or more hardware computer processors are further configured to execute the plurality of computer executable instructions in order to cause the computer system to: for the generated data item cluster: generate an alert, the alert comprising a the alert score, the one or more human-readable conclusions, the data items associated with the data item cluster, and metadata associated with the data items of the data item cluster.

According to another embodiment, the one or more hardware computer processors are further configured to execute the plurality of computer executable instructions in order to cause the computer system to: generate a user interface including a list of user-selectable alert indicators, an alert indicator being provided for each of the generated alerts, each of the alert indicators providing a summary of information associated with respective generated alerts.

According to yet another embodiment, the one or more hardware computer processors are further configured to execute the plurality of computer executable instructions in order to cause the computer system to: in response to a selection of an alert indicator by a human analyst: generate an alert display, the alert display including at least an indication of the alert score and a list of the one or more human-readable conclusions.

15

According to another embodiment, the one or more human-readable conclusions each comprise a phrase or sentence including one or more indications of summary or aggregated data associated with a plurality of the data items of the data item cluster.

According to yet another embodiment, the one or more human-readable conclusions are useable by a human analyst to perform an evaluation of the associated data item cluster.

According to another embodiment, generating the one or more human-readable conclusions comprises: selecting, based on the data cluster type associated with the particular data cluster, one or more conclusion templates; and populating the one or more conclusion templates with data associated with the particular data cluster.

According to an embodiment, a computer system comprises: one or more computer readable storage devices configured to store: a plurality of computer executable instructions; a data clustering strategy; and a plurality of data items including at least: email data items, each of the email data items including at least a subject and a sender, each of the email data items potentially associated with phishing activity; and phishing-related data items related to customers of an organization, the phishing-related data items including indicators of at least one of: customers of the organization or URLs identified as malicious by a third-party service; and one or more hardware computer processors in communication with the one or more computer readable storage devices and configured to execute the plurality of computer executable instructions in order to cause the computer system to: receive email data items from customers of the organization, each of the email data items including at least a subject and a sender; designate each of the received email data items as seeds; and for each of the designated seeds, generate a data item cluster based on the data clustering strategy by at least: adding the seed to the data item cluster; determining the subject and the sender associated with the seed; accessing the one or more computer readable storage devices and identifying one or more of the email data items having a same subject as the determined subject or a same sender as the determined sender; adding the identified one or more email data items to the data item cluster; adding one or more URLs from the email data items to the data item cluster; and in response to determining that the data item cluster includes at least a predetermined threshold quantity of email data items, designating the data item cluster.

According to another embodiment, the one or more hardware computer processors are further configured to execute the plurality of computer executable instructions in order to cause the one or more hardware computer processors to: for each designated data item cluster, initiate further automated investigation including at least: comparing URLs included in the designated data item cluster with URLs previously identified as malicious by a third-party service; and based on the comparing, identifying any URLs included in the designated data item cluster not previously identified as malicious by the third-party service; and notifying the third-party service of the identified URLs.

According to yet another embodiment, the predetermined threshold quantity is 1, 2, 3, 4, or 5.

According to another embodiment, generating the data item cluster based on the data clustering strategy further comprises: determining any new subjects or new senders associated with email data items of the data item cluster that are different from the determined subjects or senders; identifying a second one or more of the email data items having a subject that is the same as the determined new subject, or a

16

sender that is the same as the determined new sender; and adding the identified second one or more email data items to the data item cluster.

According to yet another embodiment, the identified one or more email data items are added to the data item cluster only if they were received by an initial recipient within a period of time determined based on the seed.

According to another embodiment, the period of time comprises at least one of a number of hours, a number of days, or a number of weeks, and wherein the received email data item designated as the seed was received by an initial recipient within the time period.

According to yet another embodiment, the period of time is further determined based on other email data items in the data item cluster.

According to another embodiment, generating the data item cluster based on the data clustering strategy further comprises: identifying additional one or more data items associated with any data items of the data item cluster; and adding, to the data item cluster, the additional one or more data items.

According to yet another embodiment, the one or more hardware computer processors are further configured to execute the plurality of computer executable instructions in order to cause the one or more hardware computer processors to: continuously receive email data items from customers of the organization, store the received email data items, designate the received email data items as seeds, and generate, for each of the designated seeds, data items clusters based on the data clustering strategy.

According to another embodiment, the one or more computer readable storage devices are further configured to store: a plurality of data cluster analysis rules associated with the data clustering strategy, and the one or more hardware computer processors are further configured to execute the plurality of computer executable instructions in order to cause the one or more hardware computer processors to: for each generated data item cluster: access the plurality of data cluster analysis rules associated with the data clustering strategy; analyze the data item cluster based on the accessed data cluster analysis rules; and based on the analysis of the data item cluster: determine an alert score for the data item cluster; and generate one or more human-readable conclusions regarding the data item cluster.

According to yet another embodiment, the alert score indicates a degree of correlation between characteristics of the data item cluster and the accessed data cluster analysis rules.

According to another embodiment, the degree of correlation is based on both an assessment of risk associated with the particular data cluster and a confidence level in accuracy of the assessment of risk.

According to yet another embodiment, a relatively higher alert score indicates a data cluster that is relatively more important for a human analyst to evaluate, and a relatively lower alert score indicated a data cluster that is relatively less important for the human analyst to evaluate.

According to another embodiment, each alert score for respective data clusters is assigned to a category indicating a high degree of correlation, a medium degree of correlation, or a low degree of correlation, and wherein the high degree of correlation is associated with a first color, the medium degree of correlation is associated with a second color, and the low degree of correlation is associated with a third color.

According to yet another embodiment, the one or more hardware computer processors are further configured to execute the plurality of computer executable instructions in order to cause the computer system to: for each generated data

17

item cluster: generate an alert, the alert comprising a the alert score, the one or more human-readable conclusions, the data items associated with the data item cluster, and metadata associated with the data items of the data item cluster.

According to another embodiment, the one or more hardware computer processors are further configured to execute the plurality of computer executable instructions in order to cause the computer system to: generate a user interface including a list of user-selectable alert indicators, an alert indicator being provided for each of the generated alerts, each of the alert indicators providing a summary of information associated with respective generated alerts.

According to yet another embodiment, the one or more hardware computer processors are further configured to execute the plurality of computer executable instructions in order to cause the computer system to: in response to a selection of an alert indicator by a human analyst: generate an alert display, the alert display including at least an indication of the alert score and a list of the one or more human-readable conclusions.

According to another embodiment, the one or more human-readable conclusions each comprise a phrase or sentence including one or more indications of summary or aggregated data associated with a plurality of the data items of the data item cluster.

According to yet another embodiment, the one or more human-readable conclusions are useable by a human analyst to perform an evaluation of the associated data item cluster.

According to another embodiment, generating the one or more human-readable conclusions comprises: selecting, based on the data cluster type associated with the particular data cluster, one or more conclusion templates; and populating the one or more conclusion templates with data associated with the particular data cluster.

In various embodiments, computer-implemented methods are disclosed in which, under control of one or more hardware computing devices configured with specific computer executable instructions, one or more aspects of the above-described embodiments are implemented and/or performed.

In various embodiments, a non-transitory computer-readable storage medium storing software instructions is disclosed that, in response to execution by a computer system having one or more hardware processors, configure the computer system to perform operations comprising one or more aspects of the above-described embodiments.

Advantageously, according to various embodiments, the disclosed techniques provide a more effective starting point and/or summary for an investigation of data items of various types. An analyst may be able to start an investigation from a cluster of related data items instead of an individual data item, which may reduce the amount of time and effort required to perform the investigation. The disclosed techniques may also, according to various embodiments, provide a prioritization of multiple clusters, and human-readable conclusions related to each of those clusters. For example, the analyst may also be able to start the investigation from a high priority cluster, which may allow the analyst to focus on the most important investigations, and may quickly evaluate that cluster based on the human-readable conclusions generated by the system. In each case, the processing and memory requirements of such an investigation may be significantly reduced due to the creation and use of highly efficient cluster data structures of related data items.

BRIEF DESCRIPTION OF THE DRAWINGS

The following drawings and the associated descriptions are provided to illustrate embodiments of the present disclosure

18

and do not limit the scope of the claims. Aspects and many of the attendant advantages of this disclosure will become more readily appreciated as the same become better understood by reference to the following detailed description, when taken in conjunction with the accompanying drawings, wherein:

FIG. 1 is a block diagram illustrating an example data analysis system, according to an embodiment of the present disclosure.

FIG. 2 is a block diagram illustrating an example generation of clusters by the data analysis system, according to an embodiment of the present disclosure.

FIGS. 3A-3C illustrate an example growth of a cluster of related data items, according to an embodiment of the present disclosure.

FIG. 4 illustrates an example ranking of clusters by the data analysis system, according to an embodiment of the present disclosure.

FIG. 5 illustrates an example cluster analysis user interface, according to an embodiment of the present disclosure.

FIG. 6 is a flowchart of an example method of generating clusters, according to an embodiment of the present disclosure.

FIG. 7 is a flowchart of an example method of scoring clusters, according to an embodiment of the present disclosure.

FIG. 8 illustrates components of an illustrative server computing system, according to an embodiment of the present disclosure.

FIG. 9 is a flowchart of an example generalized method of the data analysis system, according to an embodiment of the present disclosure.

Cluster Analysis

FIG. 10A is a flowchart for an example method of data cluster analysis, according to an embodiment of the present disclosure.

FIG. 10B is a flowchart of an example method of alert generation, according to an embodiment of the present disclosure.

FIG. 10C illustrates various example conclusions associated with various types of data clusters, according to various embodiments of the present disclosure.

FIGS. 11-20 illustrate example data cluster analysis user interfaces of the data analysis system, according to embodiments of the present disclosure.

FIG. 21 is a flowchart of an example method of linking related alerts or data clusters, according to an embodiment of the present disclosure.

FIG. 22 illustrates an example data cluster analysis user interface in which related alerts or data clusters are linked to one another, according to an embodiment of the present disclosure.

FIG. 23 is a flowchart of an example method of updating alerts in response to cluster regeneration, according to an embodiment of the present disclosure.

Example Application of the Data Analysis System to Malware Threat Detection

FIG. 24A is a flowchart of an example of a seed generation method of the data analysis system as applied to malware threat detection, according to an embodiment of the present disclosure.

FIG. 24B is a flowchart of an example of a clustering method of the data analysis system as applied to malware threat detection, according to an embodiment of the present disclosure.

19

FIG. 24C illustrates an example data cluster analysis user interface of the data analysis system as applied to malware threat detection, according to an embodiment of the present disclosure.

FIG. 25A is a flowchart of an example of another seed generation method of the data analysis system as applied to malware threat detection, according to an embodiment of the present disclosure.

FIG. 25B is a flowchart of an example of another clustering method of the data analysis system as applied to malware threat detection, according to an embodiment of the present disclosure.

Example Application of the Data Analysis System to Network Intrusion Detection

FIG. 26A is a flowchart of an example of a seed generation method of the data analysis system as applied to network intrusion detection, according to an embodiment of the present disclosure.

FIG. 26B is a flowchart of an example of a clustering method of the data analysis system as applied to network intrusion detection, according to an embodiment of the present disclosure.

FIG. 26C illustrates an example data cluster analysis user interface of the data analysis system as applied to network intrusion detection, according to an embodiment of the present disclosure.

Example Application of the Data Analysis System to Phishing Threat Detection

FIG. 27A is a flowchart of an example of a seed generation and clustering method of the data analysis system as applied to phishing threat detection, according to an embodiment of the present disclosure.

FIG. 27B is a flowchart of an example of another seed generation and clustering method of the data analysis system as applied to phishing threat detection, according to an embodiment of the present disclosure.

FIG. 27C illustrates an example data cluster analysis user interface of the data analysis system as applied to phishing threat detection, according to an embodiment of the present disclosure.

DETAILED DESCRIPTION

Although certain preferred embodiments and examples are disclosed below, inventive subject matter extends beyond the specifically disclosed embodiments to other alternative embodiments and/or uses and to modifications and equivalents thereof. Thus, the scope of the claims appended hereto is not limited by any of the particular embodiments described below. For example, in any method or process disclosed herein, the acts or operations of the method or process may be performed in any suitable sequence and are not necessarily limited to any particular disclosed sequence. Various operations may be described as multiple discrete operations in turn, in a manner that may be helpful in understanding certain embodiments; however, the order of description should not be construed to imply that these operations are order dependent. Additionally, the structures, systems, and/or devices described herein may be embodied as integrated components or as separate components. For purposes of comparing various embodiments, certain aspects and advantages of these embodiments are described. Not necessarily all such aspects or advantages are achieved by any particular embodiment. Thus, for example, various embodiments may be carried out in a manner that achieves or optimizes one advantage or group

20

of advantages as taught herein without necessarily achieving other aspects or advantages as may also be taught or suggested herein.

TERMS

In order to facilitate an understanding of the systems and methods discussed herein, a number of terms are defined below. The terms defined below, as well as other terms used herein, should be construed broadly to include, without limitation, the provided definitions, the ordinary and customary meanings of the terms, and/or any other implied meanings for the respective terms. Thus, the definitions below do not limit the meaning of these terms, but only provide example definitions.

Ontology: Stored information that provides a data model for storage of data in one or more databases. For example, the stored data may comprise definitions for object types and property types for data in a database, and how objects and properties may be related.

Database: A broad term for any data structure for storing and/or organizing data, including, but not limited to, relational databases (for example, Oracle database, MySQL database, and the like), spreadsheets, XML files, and text file, among others. The various terms “database,” “data store,” and “data source” may be used interchangeably in the present disclosure.

Data Item (Item), Data Object (Object), or Data Entity (Entity): A data container for information representing specific things in the world that have a number of definable properties. For example, a data item may represent an item such as a person, a place, an organization, an account, a computer, an activity, a market instrument, or other noun. A data item may represent an event that happens at a point in time or for a duration. A data item may represent a document or other unstructured data source such as an e-mail message, a news report, or a written paper or article. Each data item may be associated with a unique identifier that uniquely identifies the data item. The data item’s attributes (for example, meta-data about the data item) may be represented in one or more properties. The terms “data item,” “data object,” “data entity,” “item,” “object,” and “entity” may be used interchangeably and/or synonymously in the present disclosure.

Item (or Entity or Object) Type: Type of a data item (for example, Person, Event, or Document). Data item types may be defined by an ontology and may be modified or updated to include additional data item types. An data item definition (for example, in an ontology) may include how the data item is related to other data items, such as being a sub-data item type of another data item type (for example, an agent may be a sub-data item of a person data item type), and the properties the data item type may have.

Properties: Also referred to as “metadata,” includes attributes of a data item that represent individual data items. At a minimum, each property of a data item has a property type and a value or values. Properties/metadata associated with data items may include any information relevant to that object. For example, properties associated with a person data item may include a name (for example, John Doe), an address (for example, 123 S. Orange Street), and/or a phone number (for example, 800-0000), among other properties. In another example, metadata associated with a computer data item may include a list of users (for example, user1, user 2, and the like), and/or an IP (internet protocol) address, among other properties.

Property Type: The type of data a property is, such as a string, an integer, or a double. Property types may include

complex property types, such as a series data values associated with timed ticks (for example, a time series), and the like.

Property Value: The value associated with a property, which is of the type indicated in the property type associated with the property. A property may have multiple values.

Link: A connection between two data objects, based on, for example, a relationship, an event, and/or matching properties. Links may be directional, such as one representing a payment from person A to B, or bidirectional.

Link Set: Set of multiple links that are shared between two or more data objects.

Seed: One or more data items that may be used as a basis, or starting point, for generating a cluster. A seed may be generated, determined, and/or selected from one or more sets of data items according to a seed generation strategy. For example, seeds may be generated from data items accessed from various databases and data sources including, for example, databases maintained by financial institutions, government items, private items, public items, and/or publicly available data sources.

Cluster: A group or set of one or more related data items/objects/items. A cluster may be generated, determined, and/or selected from one or more sets of data items according to a cluster generation strategy. A cluster may further be generated, determined, and/or selected based on a seed. For example, a seed may comprise an initial data item of a cluster. Data items related to the seed may be determined and added to the cluster. Further, additional data items related to any clustered data item may also be added to the cluster iteratively as indicated by a cluster generation strategy. Data items may be related by any common and/or similar properties, metadata, types, relationships, and/or the like. Clusters may also be referred to herein as “clustered data structures,” “data item clusters,” and “data clusters.”

Seed/Cluster Generation Strategy (also referred to herein as **Seed/Cluster Generation Rule(s)**): Seed and cluster generation strategies/rules indicate processes, methods, and/or strategies for generating seeds and generating clusters, respectively. For example, a seed generation strategy may indicate that data items having a particular property (for example, data items that are credit card accounts) are to be designated as seeds. In another example, a cluster generation strategy may indicate that data items having particular properties in common with (or similar to) a seed or other data item in a cluster are to be added to the cluster. Seed and/or cluster generation strategies may specify particular searches and/or rule matches to perform on one or more sets of data items. Execution of a seed and/or cluster generation strategy may produce layers of related data items. Additionally, a seed/cluster generation strategy/rule may include multiple strategies, sub-strategies, rules, and/or sub-rules.

Overview

This disclosure relates to a data analysis system (also referred to herein as the “system”) in which memory-efficient clustered data structures (also referred to herein as “clusters”) of related data items may be automatically generated and analyzed, and results may be provided for interaction from an analyst, for example. Generation of clusters may begin by automatic generation, determination, and/or selection of an initial data item of interest, called a “seed.” As mentioned above, a data item may include any data, information, or things, such as a person, a place, an organization, an account, a computer, an activity, and event, and/or the like. Seeds may be automatically selected/generated according to various seed determination strategies, and clusters of related data items may be generated based on those seeds and according to cluster generation strategies (also referred to herein as “clus-

ter strategies,” “clustering strategies,” and/or “cluster generation rules”). Seeds and related data items may be accessed from various databases and data sources including, for example, databases maintained by financial institutions, government entities, private entities, public entities, and/or publicly available data sources. Such databases and data sources may include a variety of information and data, such as, for example, personal information, financial information, tax-related information, computer network-related data, and/or computer-related activity data, among others. Further, the databases and data sources may include various relationships that link and/or associate data items with one another. Various data items and relationships may be stored across different systems controlled by different items and/or institutions. According to various embodiments, the data analysis system may bring together data from multiple data sources in order to build clusters.

The automated analysis of the clusters may include an automated application of various criteria or rules so as to generate a compact, human-readable analysis of the data clusters. The human-readable analyses (also referred to herein as “summaries” or “conclusions”) of the data clusters may be organized into an interactive user interface so as to enable an analyst to quickly navigate among information associated with various data clusters and efficiently evaluate those data clusters in the context of, for example, a fraud investigation. Embodiments of the present disclosure also disclose automated scoring of the clustered data structures by the data analysis system. The interactive user interface may be updated based on the scoring, directing the human analyst to more critical data clusters (for example, data clusters more likely to be associated with fraud) in response to the analyst’s inputs.

In various embodiments, the data analysis system may enable an analyst (and/or other user) to efficiently perform analysis and investigations of various data clusters and related data items. For example, the system may enable an analyst to perform various financial and security investigations of data clusters of related data items. In such an investigation, the system may automatically create clusters of related data items, generate human-readable conclusions of the clusters, score those clusters, and generates an interactive user interface in which, in response to inputs from the analyst, information related to the clusters may be efficiently provided to the analyst. For example, a credit card account may be a seed that is linked by the system to various data items including, for example, customer identifiers and/or phone numbers associated with the credit card account. Further, the system may link, for example, various other credit card accounts related to the customer identifiers, to the seed credit card account. Accordingly, in various embodiments, the system may automatically cluster of various layers of data items related to the seed credit card account. One or more rules or criteria may then automatically be applied to the cluster so as to generate one or more compact, human-readable analyses (also referred to herein as “summaries” or “conclusions”) of the data clusters. The human-readable analyses may comprise phrases or sentences that provide highly relevant, and easily evaluated (by a human), information regarding the data in the cluster (for example, data items and metadata). For example, a conclusion in the current example may be “4 customer identifiers are associated with the current cluster,” or “The 2 credit card accounts in the cluster have been used in 3 different countries.” Such conclusions in an investigation may, in an embodiment, enable the analyst to determine a likelihood of fraudulent activity associated with the cluster. Further, the data items of the cluster may then be linked to possible

23

fraudulent activity. For example, the seed credit card account and the additional credit card accounts may all be linked to the potentially fraudulent activity. As mentioned above, in such an investigation the analyst may efficiently determine likely fraud, as well as discover relationships between the additional credit card accounts and the seed credit card account through several layers of related data items. Such techniques, enabled by various embodiments of the data analysis system, may be particularly valuable for investigations in which relationships between data items may include several layers, and in which such relationships may be otherwise very difficult or impossible to manually identify.

In various embodiments, the data analysis system may automatically generate, or determine, seeds based on a seed generation strategy (also referred to as “seed generation rules”). For example, for a particular set of data items, the data analysis system may automatically generate, based on a seed generation strategy, seeds by designating particular data items (and/or groups of data items) as seeds. Examples of various seed generation strategies are described below.

Further, in various embodiments, the data analysis system may automatically discover data items related to a seed, and store the resulting relationships and related data items together in a “cluster” (or, alternatively, designating the seed as the initial cluster (or initial data item of the cluster) and adding the discovered data items of the cluster). A cluster generation strategy may specify particular searches to perform at each step of an investigation, or cluster generation, process. Such searches may produce layers of related data items to add to the cluster. Further, according to an embodiment, multiple clusters may be merged and/or collapsed into a single cluster when the multiple clusters share one or more common data items and/or properties. Thus, according to an embodiment, an analyst may start an investigation with the resulting cluster, rather than the seed alone. Starting with the cluster, and associated human-readable conclusions, the analyst may form opinions regarding the related data items, conduct further analysis of the related data items, and/or may query for additional related data items.

According to various embodiments, the data analysis system may further generate various “cluster scores.” Cluster scores may include scores based on various characteristics and/or attributes associated with the cluster and/or the various data items of the cluster. In various embodiments, the data analysis system may also generate “cluster metascores” which may include, for example, an overall cluster score. Cluster metascores may, for example, be based on a combination of cluster scores of a cluster associated with a seed. In an embodiment, the system may further generate “alert scores.” Alert scores may be the same as, similar to, and/or based on any of the cluster scores, metascores, and/or conclusions described herein. In an embodiment, the alert score may be a metascore, and may be one of multiple values corresponding to, for example, a high alert, a medium alert, or a low alert. The alert score is described in further detail below. Further, cluster scores may be based on one or more generated conclusions related to the cluster, and/or the conclusions may be generated based on cluster scores.

Further, in various embodiments, for a particular set of data items, multiple clusters may be generated by the data analysis system. For example, the data analysis system may generate multiple seeds according to a seed generation strategy, and then multiple clusters based on those seeds (and based on a cluster generation strategy). In such embodiments, the data analysis system may prioritize the multiple generated clusters based upon cluster scores and/or cluster metascores. In an embodiment, the data analysis system may provide a user

24

interface including a display of human-readable conclusions of the clusters, cluster scores, cluster metascores, and/or various other cluster information. Such a user interface may be organized according to a prioritization of clusters. In various embodiments, cluster prioritization may assist an analyst in selecting particular clusters to investigate.

In various embodiments, the interactive user interface generated by the system may provide a list of clusters according to one or more alert scores (as mentioned above and described in detail below). Further, in response to an analyst selecting a cluster, information associated with the cluster may be provided to the analyst. For example, the analyst may be provided with a name of the cluster, a cluster strategy by which the cluster was generated, a list of generated conclusions, and/or one or more lists or tables of data related to the cluster. For example, the one or more lists or tables of data related to the cluster may be drawn from the data items of the cluster, and may be filtered by the analyst according to time and/or type of data. In an embodiment, various generated clusters in the interactive user interface may be organized according to clustering strategies whereby each of the clusters were generated. In an embodiment, a cluster type may be associated with each cluster, and may be determined according to the cluster strategy that generated the cluster.

As mentioned above, in various embodiments, a generated cluster may include far fewer data items than are included in a full source database, and the data items included in the cluster may only include those data items that are relevant to a particular investigation (for example, a fraud investigation). Accordingly, in an embodiment, processing of the generated cluster may be highly efficient as compared to the collection of data described above. This may be because, for example, a given fraud investigation by an analyst (for example, as the analyst sifts and/or searches through data items of a cluster) may only require storage in memory of a single cluster data structure. Further, a number of data items in a cluster may be several orders of magnitude smaller than in the entire electronic collection of data described above because only data items related to each other are included in the cluster.

Additionally, the automated analysis and scoring of clusters (as mentioned above) may enable highly efficient evaluation of the various data clusters by a human analyst. For example, the interactive user interface is generated so as to enable an analyst to quickly view critical data clusters (as determined by the automated scoring), and then in response to analyst inputs, view and interact with the generated information (including, for example, the human-readable conclusions) associated with the clusters. In response to user inputs the user interface may be updated to display raw data associated with each of the generated clusters if the analyst desires to dive deeper into data associated with a given cluster.

In various embodiments, the data analysis system may be used in various data analysis applications. Such applications may include, for example, financial fraud detection, tax fraud detection, beaconing malware detection, malware user-agent detection, other types of malware detection, activity trend detection, health insurance fraud detection, financial account fraud detection, detection of activity by networks of individuals, criminal activity detection, network intrusion detection, detection of phishing efforts, money laundering detection, and/or financial malfeasance detection. Examples of many of the above-mentioned data analysis applications, including methods and systems for identifying data items, generating data clusters, and analyzing/scoring clusters, are disclosed in the various related applications listed above and previously incorporated by reference herein.

25

Application of the data analysis system to malware threat detection, according to an embodiment, is described below in reference to FIGS. 24A-24C, 25A-25B. Application of the data analysis system to network intrusion detection, according to an embodiment, is described below in reference to FIGS. 26A-26C. Application of the data analysis system to phishing detection, according to an embodiment, is described below in reference to FIGS. 27A-27C.

In the following description, numerous specific details are set forth to provide a more thorough understanding of various embodiments of the present disclosure. However, it will be apparent to one of skill in the art that the systems and methods of the present disclosure may be practiced without one or more of these specific details.

Examples of Data Items, Properties, and Links

In various embodiments, different types of data items may have different property types. For example, a "Person" data item may have an "Eye Color" property type and an "Event" data item may have a "Date" property type. Each property as represented by data in a database may have a property type defined by an ontology used by the database. Further, data items may be instantiated in a database in accordance with a corresponding object definition for the particular data item in the ontology. For example, a specific monetary payment (for example, an item of type "event") of US\$30.00 (for example, a property of type "currency" having a property value of "US\$30.00") taking place on Mar. 27, 2009 (for example, a property of type "date" having a property value of "Mar. 27, 2009") may be stored in the database as an event object with associated currency and date properties as defined within the ontology.

Data objects defined in an ontology may support property multiplicity. In particular, a data item may be allowed to have more than one property of the same property type. For example, a "Person" data object may have multiple "Address" properties or multiple "Name" properties.

A link represents a connection between two data items and may be through any of a relationship, an event, and/or matching properties. A link may be asymmetrical or symmetrical. For example, "Person" data item A may be connected to "Person" data item B by a "Child Of" relationship (where "Person" data item B has an asymmetric "Parent Of" relationship to "Person" data item A), a "Kin Of" symmetric relationship to "Person" data item C, and an asymmetric "Member Of" relationship to "Organization" data item X. The type of relationship between two data items may vary depending on the types of the data items. For example, "Person" data item A may have an "Appears In" relationship with "Document" data item Y or have a "Participate In" relationship with "Event" data item E. As an example of an event connection, two "Person" data items may be connected by an "Airline Flight" data item representing a particular airline flight if they traveled together on that flight, or by a "Meeting" data item representing a particular meeting if they both attended that meeting. In one embodiment, when two data items are connected by an event, they are also connected by relationships, in which each data item has a specific relationship to the event, such as, for example, an "Appears In" relationship.

As an example of a matching properties connection, two "Person" data items representing a brother and a sister may both have an "Address" property that indicates where they live. If the brother and the sister live in the same home, then their "Address" properties likely contain similar, if not identical property values. In one embodiment, a link between two data item may be established based on similar or matching properties (for example, property types and/or property val-

26

ues) of the data item. These are just some examples of the types of connections that may be represented by a link and other types of connections may be represented; embodiments are not limited to any particular types of connections between data items. For example, a document may contain references to two different items. For example, a document may contain a reference to a payment (one data item), and a person (a second data item). A link between these two data items may represent a connection between these two items through their co-occurrence within the same document.

Each data item may have multiple links with another data item to form a link set. For example, two "Person" data items representing a husband and a wife may be linked through a "Spouse Of" relationship, a matching "Address" property, and/or one or more matching "Event" properties (for example, a wedding). Each link, as represented by data in a database, may have a link type defined by the database ontology used by the database.

In various embodiments, the data analysis system may access various data items and associated properties from various databases and data sources. Such databases and data sources may include a variety of information and data, such as, for example, personal information (for example, names, addresses, phone numbers, personal identifiers, and the like), financial information (for example, financial account information, transaction information, balance information, and the like), tax-related information (for example, tax return data, and the like), computer network-related data (for example, network traffic information, IP (Internet Protocol) addresses, user account information, domain information, network connection information, and the like), and/or computer-related activity data (for example, computer events, user actions, and the like), among others.

Description of the Figures

Embodiments of the disclosure will now be described with reference to the accompanying Figures, wherein like numerals refer to like elements throughout. The terminology used in the description presented herein is not intended to be interpreted in any limited or restrictive manner, simply because it is being utilized in conjunction with a detailed description of certain specific embodiments of the disclosure. Furthermore, embodiments of the disclosure described above and/or below may include several novel features, no single one of which is solely responsible for its desirable attributes or which is essential to practicing the embodiments of the disclosure herein described.

I. Example Data Analysis System

FIG. 1 is a block diagram illustrating an example data analysis system 100, according to one embodiment. As shown in the embodiment of FIG. 1, the data analysis system 100 includes an application server 115 running on a server computing system 110, a client 135 running on a client computer system 130, and at least one database 140. Further, the client 135, application server 115, and database 140 may communicate over a network 150, for example, to access data sources 160.

The application server 115 may include a cluster engine (also referred to as a "rules engine") 120, a workflow engine 125, and a user interface engine 126. The cluster engine 120, a workflow engine 125, and user interface engine 126 may be software modules as described below in reference to FIG. 8. According to an embodiment, the cluster/rules engine 120 is configured to build one or more clusters of related data items according to a defined cluster generation strategy (including generating seeds according to seed generation strategies/

27

rules), score clusters according to a scoring strategy, and/or analyze clusters including generating human-readable conclusions according to analysis rules/criteria. The cluster/rules engine 120 may read data from a variety of data sources 160 to generate seeds, generate clusters from seeds, score clusters, and analyze clusters. Once created, the resulting clusters may be stored on the server computing system 110 and/or on the database 140. The operations of the cluster/rules engine 120 are discussed in detail below.

As mentioned, in an embodiment, the cluster/rules engine 120 is configured to score the clusters, according to a defined scoring strategy. The score may indicate the importance of analyzing the cluster. For instance, the cluster/rules engine 120 may execute a scoring strategy that aggregates the account balances of credit card accounts within the cluster. Because, for example, a large aggregated total balance may indicate a large liability for a financial institution, a cluster with such a large total balance may be considered to have a higher score relative to other clusters with lower aggregated total balances (and, therefore, lower scores). Thus, a cluster with a higher score relative to a cluster with a lower score may be considered more important to analyze.

As described below, in an embodiment the cluster/rules engine 120 is configured to apply one or more analysis rules or criteria to the generated cluster to generate one or more human-readable conclusions (as mentioned above, also referred to herein as “summaries”). In various embodiments the one or more analysis rules/criteria may be based on one or more scoring strategies. Also, in various embodiments the scoring strategies may be based on one or more analysis rules/criteria. As described below, the cluster/rules engine 120 may generate an “alert score” for a given cluster. The alert score may be the same as, similar to, and/or based on any of the cluster scores, metascopes, and/or conclusions described herein. In an embodiment, the alert score may be a metascore, and may be one of multiple values corresponding to, for example, a high alert, a medium alert, or a low alert. The alert score is described in further detail below.

In an embodiment, the user interface engine 126 generates various user interfaces of the data analysis system as described below. In one embodiment, the cluster engine 120, in conjunction with the user interface engine 126, organizes and presents the clusters according to the assigned scores. The cluster engine 120 and the user interface engine 126 may present information associated with the clusters and/or interactive representations of the clusters within a user interface presented to the analyst, as described below. For example, the representations may provide visual indications (e.g., graphs or other visualizations) of the related data items within the clusters. The cluster engine 120 and/or the user interface engine 126 may generate the user interface, such as a web application and/or a dynamic web page displayed within the client 135. The cluster engine 120 and/or the user interface engine 126 may also allow an analyst to create tasks associated with the clusters. Example operations of the cluster engine 120 and/or the user interface engine 126 are discussed in detail below in conjunction with various figures. In one embodiment, the cluster engine 120 generates clusters automatically, for example, for subsequent review by analysts.

Analysts may also assign tasks to themselves or one another via a workflow user interface generated by the workflow engine 125 and/or the user interface engine 126, for example. The workflow engine 125 and/or the user interface engine 126 may consume scores generated by the cluster engine 120. For example, the workflow engine 125 and/or the user interface engine 126 may present an analyst with clusters generated, scored, and ordered by the cluster engine 120.

28

The client 135 may represent one or more software applications or modules configured to present data and translate input, from the analyst, into requests for data analyses by the application server 115. In one embodiment, the client 135 and the application server 115 may be embodied in the same software module and/or may be included in the same computing system. However, several clients 135 may execute on the client computer 130, and/or several clients 135 on several client computers 130 may interact with the application server 115. In one embodiment, the client 135 may be a browser accessing a web service.

While the client 135 and application server 115 are shown running on distinct computing systems, the client 135 and application server 115 may run on the same computing system. Further, the cluster engine 120 and the workflow engine 125 may run on separate applications servers 115, on separate server computing systems, or some combination thereof. Additionally, a history service may store the results generated by an analyst relative to a given cluster.

In one embodiment, the data sources 160 provide data available to the cluster engine to create or generate seeds and/or to create or generate clusters from a seed or a set of seeds. Such data sources may include relational data sources, web services data, XML data, and the like. Further, such data sources may include a variety of information and data, for example, personal information, financial information, tax-related information, computer network-related data, and/or computer-related activity data, among others. For example, the data sources may be related to customer account records stored by a financial institution. In such a case, the data sources may include a credit card account data, bank account data, customer data, and transaction data. The data may include data attributes such as account numbers, account balances, phone numbers, addresses, and transaction amounts, and the like. Of course, data sources 160 is included to be representative of a variety of data available to the server computer system 110 over network 150, as well as locally available data sources.

The database 140 may be a Relational Database Management System (RDBMS) that stores the data as rows in relational tables. The term “database,” as used herein, may refer to an database (e.g., RDBMS or SQL database), or may refer to any other data structure, such as, for example a comma separated values (CSV), extensible markup language (XML), text (TXT) file, flat file, spreadsheet file, and/or any other widely used or proprietary format. While the database 140 is shown as a distinct computing system, the database 140 may operate on the same server computing system 110 as the application server 115.

II. Example Cluster Generation

FIG. 2 is a block diagram illustrating an example generation of clusters by data analysis system 200, according to an embodiment. As shown, in an embodiment the cluster engine 120 (FIG. 1) interacts with a seed list 210, a cluster list 250, a cluster strategy store 230, and data bindings 237. The seed list 210 may include seeds 212-1, 212-2 . . . 212-S, and the cluster list 250 may include clusters 252-1, 252-2 . . . 252-C. The cluster engine 120 may be configured as a software application, module, or thread that generates the clusters 252-1, 252-2 . . . 252-C from the seeds 212-1, 212-2 . . . 212-S.

Seeds 212 (including one, some, or all of seeds 212-1 through 212-S) may be generated by the cluster engine 120 according to various seed generation strategies/rules. Examples of seed generation are described below in reference to various example applications of the data analysis system.

According to an embodiment, once generated, seeds **212** may be the starting point for generating a cluster **252**. To generate a cluster, the cluster engine **120** may retrieve a given seed **212** from the seed list **210**. The seed **212** may be a data item or group of data items within the database **140**, such as a customer name, a customer social security number, an account number, and/or a customer telephone number.

The cluster engine **120** may generate the cluster **252** from the seed **212**. In one embodiment, the cluster engine **120** generates the cluster **252** as a collection of data items and the relationships between the various data items. As noted above, the cluster strategy may execute data bindings in order to add each additional layer of data items to the cluster. For example, the cluster engine **120** may generate the cluster **252-1** from a seed credit card account. The cluster engine **120** may first add the credit card account to the cluster **252-1**. The cluster engine **120** may then add customers related to the credit card account to the cluster **252-1**. The cluster engine **120** may complete the cluster **252-1** by adding additional credit card accounts related to those customers. As the cluster engine **120** generates the cluster **252-1**, the cluster engine **120** may store the cluster **252-1** within the cluster list **250**. The cluster **252-1** may be stored as a graph data structure or other appropriate data structure.

The cluster list **250** may be a collection of tables in the database **140**. In such a case, there may be a table for the data items of each cluster **252**, such as those of example cluster **252-1** discussed above, a table for the relationships between the various data items, a table for the attributes of the data items, and a table for scores of the clusters. The cluster list **250** may include clusters **252** from multiple investigations. Note that the cluster engine **120** may store portions of clusters **252** in the cluster list **250** as the cluster engine **120** generates the clusters **252**. Persons skilled in the art will recognize that many technically feasible techniques exist for creating and storing data structures that may be used to implement the systems and methods of the data analysis system.

The cluster strategy store **230** may include cluster strategies **232-1**, **232-2** . . . **232-N**. Each cluster strategy may include data binding references **237** to one or more data bindings **237**. As noted, each data binding may be used to identify data that may grow a cluster (as determined by the given search strategy **232**). For example, the cluster engine **120** may execute a cluster strategy **232-1** to generate the cluster **252-1**. Specifically, the cluster engine **120** may execute the cluster strategy **232-1** in response to selection of that cluster strategy by an analyst. The analyst may submit a selection of one or more cluster strategies to perform on a seed or group of seeds to the cluster engine **120** through the client **135**. Alternatively, the cluster engine **120** may automatically select one or more cluster strategies, such as based on user preferences or rules.

According to an embodiment, each cluster strategy **232** is configured so as to perform an investigation processes for generating a cluster **252**. Again, for example, the cluster strategy **232-2** may include data binding references **237** to a collection of data bindings executed to add layer after layer of data to a cluster. The investigation process may include searches to retrieve data items related to a seed **212** that is selected for clustering using cluster strategy **232-2**. For example, the cluster strategy **232-2** may start with a possibly fraudulent credit card account as the seed **212-2**. The cluster strategy **232-2** may search for customers related to the credit card account, and then additional credit card accounts related to those customers. A different cluster strategy **232-3** may search for customers related to the credit card account, phone numbers related to the customers, additional customers

related to the phone numbers, and additional credit card accounts related to the additional customers, for example.

In an embodiment, cluster strategies **232** include references to at least one data binding **237** (such as data bindings **237-1** through **237-3**). The cluster engine **120** may execute a search protocol specified by the data binding **237** to retrieve data, and the data returned by a given data binding may form a layer within the cluster **252**. For instance, the data binding **237** (and/or the search protocol of the data binding **237**) may retrieve sets of customers related to an account by an account owner attribute. The data binding **237** (and/or the search protocol of the data binding **237**) may retrieve the set of related data items from a data source. For instance, the data binding **237-1** may specify a database query to perform against a database. Likewise, the data binding **237-2** may define a connection and/or query to a remote relational database system and the data binding **237-3** may define a connection and/or query against a third-party web service. Once retrieved, the cluster strategy **232** may evaluate whether the returned data should be added to a cluster being grown from a given seed **212**.

Multiple cluster strategies **232** may reference a given data binding **237**. The analyst may update the data binding **237**, but typically updates the data binding **237** only if the associated data source changes. A cluster strategy **232** may also include a given data binding **237** multiple times. For example, executing a data binding **237** using one seed **212** may generate additional seeds for that data binding **237** (and/or generate seeds for another data binding **237**). More generally, different cluster strategies **232-1**, **232-2** . . . **232-N** may include different arrangements of various data bindings **237** to generate different types of clusters **252**.

The cluster strategies **232** may specify that the cluster engine **120** use an attribute from the related data items retrieved with one data binding **237**, as input to a subsequent data binding **237**. The cluster engine **120** may use the subsequent data binding **237** to retrieve a subsequent layer of related data items for the cluster **252**. For instance, a particular cluster strategy **232** may specify that the cluster engine **120** retrieve a set of credit card account data items with a first data binding **237-1**. That cluster strategy **232** may also specify that the cluster engine **120** then use the account number attribute from credit card account data items as input to a subsequent data binding **237-2**. The cluster strategy **232** may also specify filters for the cluster engine **120** to apply to the attributes before performing the subsequent data binding **237**. For instance, if the first data binding **237-1** were to retrieve a set of credit card account data items that included both personal and business credit card accounts, then the cluster engine **120** could filter out the business credit card accounts before performing the subsequent data binding **237-2**.

In operation, according to an embodiment, the cluster engine **120** generates a cluster **252-1** from a seed **212-1** by first retrieving a cluster strategy **232**. Assuming the analyst selected a cluster strategy **232-2**, the cluster engine **120** would retrieve the cluster strategy **232-2** from the cluster strategy store **230**. The cluster engine **120** may then retrieve the seed **212-1** as input to the cluster strategy **232-2**. The cluster engine **120** may execute the cluster strategy **232-2** by retrieving sets of data by executing data bindings **237** referenced by the cluster strategy **232-2**.

For example, the cluster strategy **232-2** may execute data bindings **237-1**, **237-2**, and **237-3**. Accordingly, the cluster engine **120** may evaluate data returned by each data binding **237** to determine whether to use that data to grow the cluster **252-1**. The cluster engine **120** may then use elements of the returned data as input to the next data binding **237**. Of course,

31

a variety of execution paths are possible for the data bindings 237. For example, assume one data binding 237 returned a set of phone numbers. In such a case, another data binding 237 may evaluate each phone number individually. As another example, one data binding 237 may use input parameters obtained by executing multiple, other data bindings 237. More generally, the cluster engine 120 may retrieve data for each data binding referenced by the cluster strategy 232-2. The cluster engine 120 may then store the complete cluster 252-1 in the cluster list 250.

As the cluster engine 120 generates the clusters 252-1, 252-2 . . . 252-C from seeds 212-1, 212-2 . . . 212-S, the cluster list 250 may include overlapping clusters 252. For example, two clusters 252-1 and 252-C may overlap if both clusters 252-1 and 252-C include a common data item. In an example, a larger cluster 252 formed by merging two smaller clusters 252-1 and 252-C may be a better investigation starting point than the smaller clusters 252-1 and 252-C individually. The larger cluster 252 may provide additional insight or relationships, which may not be available if the two clusters 252-1 and 252-C remain separate.

In an embodiment, the cluster engine 120 includes a resolver 226 that is configured to detect and merge two or more overlapping clusters 252 together. For example, the resolver 226 may compare the data items within a cluster 252-1 to the data items within each one of the other clusters 252-2 through 252-C. If the resolver 226 finds the same data item within the cluster 252-1 and a second cluster 252-C, then the resolver 226 may merge the two clusters 252-1 and 252-C into a single larger cluster 252. For example, the cluster 252-1 and cluster 252-C may both include the same customer. The resolver 226 may compare the data items of cluster 252-1 to the data items of cluster 252-C and detect the same customer in both clusters 252. Upon detecting the same customer in both clusters 252, the resolver 226 may merge the cluster 252-1 with cluster 252-C. The resolver 226 may test each pair of clusters 252 to identify overlapping clusters 252. Although the larger clusters 252 may be better investigation starting points, an analyst may want to understand how the resolver 226 formed the larger clusters 252. Accordingly, the resolver 226, may store a history of each merge.

In various embodiments, clusters may be merged based on various criteria and/or combinations of criteria include, for example, when the clusters include a minimum number of data items that are common among the clusters, when the clusters include a minimum number of data items that are common among the clusters and which data items are within a particular proximity in each cluster to a seed of the cluster, when a particular quantity of properties are common among data items of the clusters even when the data items themselves are not identical, and/or the like.

In an embodiment, cluster merging (for example, by resolver 226) may be optionally disabled for particular types of data items, and/or particular data items. For example, when a particular data item, or type of data item, is so common that it may be included in many different clusters (for example, an institutional item such as a bank), merging of cluster based on that common item (for example, the particular bank) or common type of item (for example, banks in general) may be disabled. In another embodiment, cluster may be merged only when they share two or more common data items and/or other properties. In an embodiment, when two clusters are determined to share a data item that this very common (such that they cluster may not be merged based on that item) the system may automatically determine whether the two clusters share one or more other data items and/or properties such that they may be merged. In various embodiments, cluster merging

32

may be disabled based on other criteria. For example, cluster merging between two related clusters may be disabled when one or both of the two clusters reach a particular size (for example, include a particular number of data items).

After the cluster engine generates a group of clusters from a given collection of seeds (and after merging or resolving the cluster), the cluster engine 120 may score, rank, and/or otherwise order the clusters relative to a scoring strategy 442. In some embodiments, clusters are scored and provided to the analysis without resolving.

In one embodiment, the analysis system 100, and more specifically, the cluster engine 120, receives a request for cluster generation. In response to the request, a list of seeds may be generated, clusters may be generated based on those seeds, and the clusters may be ranked, ordered, and presented to analysts. In an embodiment, the cluster engine 120 may consume seeds generated by other systems. Alternatively, in other embodiments, cluster engine 120 may generate the seeds 212-1, 212-2 . . . 212-S. For instance, the cluster engine 120 may include a seed generation strategy (also referred to as a "lead generation strategy") that identifies data items, or groups of data items, as potential seeds 212. The seed generation (and/or lead generation) strategy may apply to a particular business type, such as credit cards, stock trading, or insurance claims, and may be run against a cluster data source 160 or an external source of information.

In an embodiment, the analysis system 100 may not include data bindings as described above. Rather, according to an embodiment, the analysis system 100 may include one or more interfaces and/or connections to various internal and/or external data stores of data items and/or other information (for example, data sources(s) 160. According to an embodiment, the system may include a generic interface and/or connection to various internal and/or external data stores of data items and/or other information. For example, the analysis system 100 may include a generic data interface through which the system may search, access, and/or filter various data item information during seed generation, cluster generation, and/or analysis of the clusters. The generic interface may include various aspects that enable searching, accessing, and/or filtering of data. For example, the generic interface may access various data sources that each have differing data formats. The generic interface may accordingly covert and/or filter the accessed data to a common format. Alternatively, the data sources may include functionality through which stored data may be searched and/or converted to a standard format automatically. In an embodiment, the generic interface may enable Federated search of multiple data stores of data item-related information. Accordingly, in various embodiments, the analysis system 100 may access various data sources for data item clustering and seed generation.

Additional details of the server computing system 110, the data sources 160, and other components of the data analysis system are described below in reference to FIG. 8.

FIGS. 3A-3C illustrate an example growth of a cluster 252 of related data items, according to an embodiment. As shown in FIG. 3A, an example cluster 252 may include a seed item 302, links 303-1 and 303-2, and related data items 305-1 and 305-2. The cluster 252 may be based upon a seed 212 (for example, data item 302). The cluster engine 120 may build the cluster 252 by executing a cluster strategy 232 with the following searches:

- Find seed owner
- Find all phone numbers related to the seed owner
- Find all customers related to the phone numbers
- Find all accounts related to the customers
- Find all new customers related to the new accounts

33

In the example, assuming the seed **212** is fraudulent credit card account, the cluster engine **120** would add the credit card account to the cluster **252** as the seed item **302**. The cluster engine **120** may then use the account owner attribute of the credit card account as input to a data binding **237**. The cluster engine **120** may execute the search protocol of the data binding **237** to retrieve the customer data identifying the owner of the fraudulent credit card account. The cluster engine **120** would then add the customer data to the cluster **252** as the related data item **305-1**. The cluster engine **120** would also add the account owner attribute as the link **303-1** that relates the account number to the customer data of the owner. The cluster engine **120** would execute the next search of the cluster strategy **232** by inputting the customer identifier attribute of the customer data into a data binding **237** to retrieve a phone data. The cluster engine **120** would then add the phone data as the related data item **305-2** and the customer identifier attribute as the link **303-2** between the customer data and the phone data. At this point in the investigation process, the cluster **252** would include the seed item **302**, two links **303-1** and **303-2**, and two related data items **305-1** and **305-2**. That is, the cluster **252** would include the fraudulent credit card account, the customer data of the owner of the credit card, and the phone number of the owner. By carrying the investigation process further, the cluster engine **120** may reveal further related information, for example, additional customers and/or potentially fraudulent credit card accounts.

Turning to FIG. 3B, and continuing the example, the cluster engine **120** may continue executing the cluster strategy **232** by searching for additional account data items related to the phone number of the owner of the fraudulent credit card account. As discussed, the phone number may be stored as related data item **305-2**. The cluster engine **120** would input the phone owner attribute of the phone number to a data binding **237**. The cluster engine **120** would execute the search protocol of data binding **237** to retrieve the data of two additional customers, which the cluster engine **120** would store as related data items **305-3** and **305-4**. The cluster engine **120** would add the phone owner attribute as the links **303-3** and **304-4** between the additional customers and the phone number.

Continuing the example, FIG. 3C shows the cluster **252** after the cluster engine **120** performs the last step of the example cluster strategy **232**. For example, the cluster engine **120** would use the customer identifier attribute of the related data item **305-3** and **305-4** to retrieve and add additional account data items as the related data items **305-5** and **305-6**. The cluster engine **120** would couple the related data items **305-5** and **305-6** to the related data items **305-3** and **305-4** with the customer identifier attributes stored as links **303-5** and **303-6**. Thus, the cluster **252** would include six related data items **305** related by six links **303**, in addition to the seed item **302**.

In an embodiment, the analyst may identify and determine whether the additional data account items, stored as related data items **305-5** and **305-6**, represent fraudulent credit card accounts more efficiently than if the analyst started an investigation with only the seed **302**. As the foregoing example illustrates, according to various embodiments, the data analysis system may enable an analyst to advantageously start an investigation with a cluster including many related data items (such as the example cluster **252** with the seed item **302** and related data items **305**) rather than a single data item.

In various embodiments, clusters may be generated automatically, on a schedule, on demand, and/or as needed, as described below.

34

III. Example Cluster Scoring/Ranking

FIG. 4 illustrates an example ranking of clusters **252** by the data analysis system **100** shown in FIG. 1, according to an embodiment of the present disclosure. As shown, an example system **400** of FIG. 4 illustrates some of the same elements as shown in FIG. 1 and FIG. 2, including the cluster engine **120** in communication with the cluster list **250**. In addition, FIG. 4 illustrates a scoring strategy store **440** in communication with the cluster engine **120**. The scoring strategy store **440** includes scoring strategies **442-1**, **442-2** . . . **442-R**.

In an embodiment, the cluster engine **120** executes a scoring strategy **442** to score a cluster **252**. For example, the cluster engine **120** may generate a cluster (for example, via a cluster strategy/data bindings) and attempt to resolve it with existing clusters. Thereafter, the cluster engine **120** may score the resulting cluster with any scoring strategies associated with a given cluster generation strategy. In an embodiment, the multiple scores may be generated for a given cluster. The multiple scores may be based on various aspects, metrics, or data associated with the cluster. In one embodiment, a cluster metascore may be generated based on a combination or aggregation of scores associated with a given cluster. Ordering for a group of clusters, (according to a given scoring strategy) may be performed on demand when requested by a client. Alternatively, the analyst may select a scoring strategy **442** through the client **135** and/or the analyst may include the selection within a script or configuration file. In another alternative, the data analysis system may automatically select a scoring strategy. In other embodiments, the cluster engine **120** may execute several scoring strategies **442** to determine a combined score for the cluster **252**.

In an embodiment, a scoring strategy (such as scoring strategy **442**) specifies an approach for scoring a cluster (such as cluster **252**). A score may indicate a relative importance or significance of a given cluster. For example, the cluster engine **120** may execute a scoring strategy **442-1** to determine a score by counting the number of a particular data item type that are included within the cluster **252**. Assume, for example, a data item corresponds with a credit account. In such a case, a cluster with a large number of accounts opened by a single individual (possibly within a short time) might correlate with a higher fraud risk. Of course, a cluster score may be related to a high risk of fraud based on the other data in the cluster, as appropriate for a given case. More generally, each scoring strategy **442** may be tailored based on the data in clusters created by a given cluster strategy **230** and a particular type of risk or fraud (and/or amounts at risk) of interest to an analyst.

According to an embodiment, the cluster engine **120** scores a cluster **252-1** by first retrieving a scoring strategy **442**. For example, assume an analyst selects scoring strategy **442-1**. In response, the cluster engine **120** may retrieve the scoring strategy **442-1**. The cluster engine **120** may also retrieve the cluster **252-1** from the cluster list **250**. After determining the score of the cluster **252-1**, the cluster engine **120** may store the score with the cluster **252-1** in the cluster list **250**.

The cluster engine **120** may score multiple clusters **252-1**, **252-2** . . . **252-C** in the cluster list **250**. The cluster engine **120** may also rank the clusters **252-1**, **252-2** . . . **252-C** based upon the scores. For instance, the cluster engine **120** may rank the cluster **252-1**, **252-2** . . . **252-C** from highest score to lowest score. In various embodiment, cluster may be ranked according to multiple scores, combinations of scores, and/or metascores.

As mentioned above, the cluster/rules engine **120** may generate an "alert score" for the clusters. The alert score may be the same as, similar to, and/or based on any of the cluster

35

scores, metascores, and/or conclusions described herein. In an embodiment, the alert score may be a metascore, and may be one of multiple values corresponding to, for example, a high alert, a medium alert, or a low alert. The alert score is described in further detail below.

IV. Example User Interface

FIG. 5 illustrates an example user interface 500, according to one embodiment. As described above, the cluster engine 120, the workflow engine 125, and/or the user interface engine 126 may be configured to present the user interface 500. As shown, the example user interface 500 includes a selection box 510, a cluster strategy box 530, a cluster summary list 525, a cluster search box 520, and a cluster review window 515. The user interface 500 may be generated as a web application or a dynamic web page displayed within the client 135.

In the example user interface 500 of FIG. 5, the selection box 510 may allow the analyst to select, for example, a seed generation strategy and/or a previously generated seed or seed list (for example, seed list 210). The analyst may select the items (for example, a seed generation strategy) by, for example, entering a name of a particular item into a dropdown box (and/or other interface element) in the selection box 510 (for example, the dropdown box showing a selected strategy “Strategy-A”) and selecting a “Go” button (and/or other interface element). Alternatively, the analyst may select a particular item by, for example, expanding the dropdown box and selecting an item from the expanded dropdown box, which may list various seed generation strategies and/or seed lists, for example. In various examples, seed lists and/or seed generation strategies may be selected by the analyst that correspond to likely fraudulent financial accounts, credit card account originating at a particular bank branch, savings accounts with balances above a particular amount, and/or any of the other seed generation strategies described below in reference to the various applications of the system.

For example, when the analyst selects a particular seed generation strategy, the system may generate a seed list (for example, seed list 210) and then may generate clusters based on seeds of the seed list. The seed list and/or clusters may, in an embodiment, be generated in response to a selection of a particular seed generation strategy. The seed generation strategy may generate a seed list (for example, seed list 210) and/or clusters (for example, clusters 252-1, 252-2, . . . 252-C of the cluster list 250) from the database 140 and/or an external source of information (for example, a cluster data source 160). Alternatively, when the analyst selects a previously generated seed or seed list (for example, seed list 210), the system may retrieve data related to the selected seed list (for example, the seed items, clusters, and/or related clustered data items) from, for example, database 140 and/or an external source of information (for example, a cluster data source 160). In an embodiment, clusters may be generated in response to a selection of a previously generated seed list (or, alternatively, a previously generated seed). Alternatively, cluster may be been previously generated, and may be retrieved in response to selection of a previously generated seed list (or, alternatively, a previously generated seed). In an embodiment, the analyst may select a particular cluster of interest via the selection box 510.

Further, in the example user interface 500 the cluster strategy box 530 displays the cluster strategies 232 that the cluster engine 120 ran against the seed list 210. The cluster engine 120 may execute multiple cluster strategies 232 against the seed list 210, so there may be multiple cluster strategies 232

36

listed in the cluster strategy box 530. The analyst may click on the name of a given cluster strategy 232 in the cluster strategy box 530 to review the clusters 252 that the cluster strategy 232 generated.

In an embodiment, the user interface 500 displays information associated with the clusters 252 in the cluster summary list 525. For example, the information associated with the clusters may include characteristics of the clusters 252, such as identifiers, scores, and/or analysts assigned to analyze the clusters 252. The system may select the clusters 252 for display in the cluster summary list 525 according to those or other characteristics. For instance, the system may display the cluster information in the order of the scores of the clusters 252, where a summary of the highest scoring cluster 252 is displayed first.

The system (for example, cluster engine 120, the workflow engine 125, and/or the user interface engine 126) may control the order and selection of the cluster information within the cluster summary list 525 based upon an input from the analyst. The cluster search box 520 may include a search text box coupled to a search button and a pull-down control. The analyst may enter a characteristic of a cluster 252 in the search text box and then instruct the workflow engine 125 to search for and display clusters 252 that include the characteristic by pressing the search button. For example, the analyst may search for clusters with a particular score. The pull-down control may include a list of different characteristics of the clusters 252, such as score, size, assigned analyst, and/or date created. The analyst may select one of the characteristics to instruct the workflow engine 125 to present the information associated with the clusters 252 arranged by that characteristic.

In an embodiment, the system is also configured to present details of a given cluster 252 within the cluster review window 515. The system displays the details of the cluster 252, for example, the score, and/or average account balances within a cluster, when the analyst clicks a mouse pointer on the associated summary within the cluster summary list 525. The system may present details of the cluster 252, such as the name of an analyst assigned to analyze the cluster 252, a score of the cluster 252, and/or statistics or graphs generated from the cluster 252. These details may allow the analyst to determine whether to investigate the cluster 252 further. The cluster review window 515 may also include a button which may be clicked to investigate a cluster 252 within a graph, and an assign button for assigning a cluster to an analyst.

An analyst may click a mouse pointer on an “Investigate in Graph” button representing a cluster to investigate the cluster within an interactive graph. The interactive representation may be a visual graph of the cluster 252, where icons represent the items of the cluster 252 and lines between the icons represent the links between items of the cluster 252. For example, the workflow engine 125 may display the interactive graph of the cluster 252 similar to the representation of the cluster 252 in FIG. 3C. The interactive representation may allow the analyst to review the attributes of the related data items and/or perform queries for additional related data items.

In an embodiment, an administrative user may click a mouse pointer on an assign button to assign the associated cluster 252 to an analyst. The workflow engine 125 may also allow the administrative user to create tasks associated with the clusters 252, while the administrative user assigns the cluster 252. For example, the administrative user may create a task for searching within the three highest scoring clusters 252 for fraudulent credit card accounts. The system may display the cluster information in the cluster summary list 525 according to the names of the analysts assigned to the clusters

37

252. Likewise, the system may only display cluster information for the subset of the clusters 252 assigned to an analyst.

The interface shown in FIG. 5 is included to illustrate one example interface useful for navigating and reviewing clusters generated using the cluster engine 120 and the workflow engine 125. In other embodiments, other user interface constructs may be used to allow the analyst to select cluster strategies 232, scoring strategies 242, and/or seed generation strategies, initiate an investigation, and/or review and analyze the clusters 252. For example, the user interface engine 126 may display additional controls within the user interface 500 for controlling the cluster generation process and selecting seed generation strategies, cluster strategies 232, and/or scoring strategies 242. Also, the user interface 500 may be displayed without the selection box 510 or the options to select a seed generation strategy. In addition, although the workflow engine 125 may generate the user interface 500, in various embodiments the user interface 500 may be generated by a software application distinct from the workflow engine 125. Further, in various embodiments, the cluster review window 515 may be configured to display a preview of the cluster 252 and/or additional statistics generated from the cluster 252. As such, an interactive representation of the cluster 252 may be presented in an additional user interface and/or the cluster 252 may be exported to another software application for review by the analyst.

In an alternative embodiment, and as described below in reference to the various figures, various other user interfaces may be generated by the system.

V. Example Operations

FIG. 6 is a flowchart of an example method of generating clusters, according to an embodiment. Although the method is described in conjunction with the systems of FIGS. 1 and 2, persons skilled in the art will understand that any system configured to perform the method, in any order, is within the scope of this disclosure. Further, the method 600 may be performed in conjunction with method 700 for scoring a cluster, described below, and the various other methods described below including analyzing a cluster.

As shown, example cluster generation method 600 begins at block 605, where the cluster engine 120 retrieves a cluster strategy (e.g., cluster strategy 232-2) and a seed 212. Once a cluster strategy is selected, the cluster engine 120 may identify a list of seeds from which to build clusters using the selected cluster strategy. At block 610, the cluster engine 120 initializes a cluster 252 with one of the seeds in the list. The cluster 252 may be stored as a graph data structure. The cluster engine 120 may initialize the graph data structure and then add the seed 212-1 to the graph data structure as the first data item.

At block 615, the cluster engine 120 may grow the cluster 252 by executing the search protocol of a data binding 237 from the cluster strategy 232-2. The cluster strategy 232-2 may include a series of data bindings 237 that the cluster engine 120 executes to retrieve related data items. A given data binding 237 may include queries to execute against a cluster data source 160 using the seed as an input parameter. For example, if the seed 212-1 is an account number, then the data binding 237 may retrieve the data identifying the owner of the account with the account number. After retrieving this information, the cluster engine 120 may add the customer data item to the cluster as a related data item and the account owner attribute as the link between the seed 212-1 and the related data item. After retrieving the related data items, the cluster engine 120 may add them to the cluster 252.

38

At block 620, the cluster engine 120 determines if the cluster strategy 232-2 is fully executed. If not the method 600 returns to block 615 to execute additional data bindings for a given seed. Alternatively, as described above, the cluster engine 120 may grow the cluster by searching for, accessing, and/or filtering various data items through, for example, a generic interface to various internal and/or external data sources. Further, in an embodiment, the cluster engine 120 may determine whether the cluster being generated is to be merged with another cluster, as described above. Once the cluster strategy is executed for that seed, the cluster engine 120 may determine and assign a score (or, alternatively, multiple scores) to that cluster (relative to a specified scoring strategy). After generating clusters for a group of seeds, such clusters may be ordered or ranked based on the relative scores. Doing so may allow an analyst to rapidly identify and evaluate clusters determined to represent, for example, a high risk of fraud.

At block 625, the cluster engine 120 may store the cluster 252 in cluster list 250. As mentioned above, the cluster list 250 may be a collection of tables within a relational database, where a table may include the seed and related data items of the cluster 252 and another table may include links between the related data items of the cluster 252.

At block 630, the cluster engine 120 determines if there are more seeds 212 to analyze in the seed list 210. If so, the method 600 returns to block 605 to generate another cluster from the next seed. Otherwise, the method 600 ends. Note, while method 600 describes a single cluster being generated, one of skill in the art will recognize that multiple instances of the cluster generation process illustrated by method 600 may be performed in parallel.

FIG. 7 is a flowchart of an example method of scoring clusters, according to an embodiment. Although the method is described in conjunction with the systems of FIGS. 1 and 4, persons skilled in the art will understand that any system configured to perform the method steps, in any order, is within the scope of the present invention.

As shown, the example cluster scoring method 700 begins at block 705, where the cluster engine 120 retrieves a scoring strategy 442 and a cluster 252 (for example, a cluster just created using the method 600 of FIG. 6). In other cases, the cluster engine 120 may retrieve the scoring strategy 442 associated with a stored cluster. Other alternatives include an analyst selecting a scoring strategy 442 through the client 135, the cluster engine 120 via the cluster analysis UI 500, a script, or a configuration file. The cluster engine 120 may retrieve the selected scoring strategy 442 from the scoring strategy store 440, and the cluster 252 from the cluster list 250.

At block 710, the cluster engine 120 executes the scoring strategy 442 against the cluster 252. The scoring strategy 442 may specify characteristics of the related data items within the cluster 252 to aggregate. The cluster engine 120 may execute the scoring strategy 442 by aggregating the specified characteristics together to determine a score. For instance, the cluster engine 120 may aggregate account balances of related data items that are account data items. In such a case, a total amount of dollars (and/or average dollars or any other aggregated, averaged, or normal attribute of the cluster) included within the balances of the account data items of the cluster 252 may be the score of the cluster 252.

At block 715, the cluster engine 120 may store the score with the cluster 252 in the cluster list 250. At step 720, the cluster engine 120 determines if there are more clusters 252 to score. For example, in one embodiment, a set of clusters may be re-scored using an updated scoring strategy. In other cases,

the cluster engine may score each cluster when it is created from a seed (based on a given cluster generation and corresponding scoring strategy). If more clusters remain to be scored (and/or re-scored), the method 700 returns to block 705.

At block 725, the cluster engine 120 may rank the clusters 252 according to the scores of the clusters 252. For example, after re-scoring a set of clusters (or, alternatively, after scoring a group of clusters generated from a set of seeds), the cluster engine 125 may rank the clusters 252 from highest score to lowest score. The ranking may be used to order a display of information associated with the clusters 252 presented to the analyst. The analyst may rely upon the ranking and scores to determine which clusters 252 to analyze first. The ranking and sorting may generally be performed on-demand when an analyst is looking for a cluster to investigate. Thus, the ranking need not happen at the same time as scoring. Further, the clusters may be scored (and later ranked) using different ranking strategies.

In various embodiments, multiple scores for each cluster may be determined according to methods similar to the example method 700. Accordingly, clusters may be ranked according to any of multiple scores. Additionally, in various embodiments, multiple scores may be combined and/or aggregated into a metascore that may be used to rank the clusters. Various example score and metascore determinations are described below in reference to FIGS. 10C, 11C, 12C, and 13C.

VI. Example Implementation Mechanisms/Systems

FIG. 8 illustrates components of an illustrative server computing system 110, according to an embodiment. The server computing system 110 may comprise one or more computing devices that may perform a variety of tasks to implement the various operations of the data analysis system. As shown, the server computing system 110 may include, one or more central processing unit (CPU) 860, a network interface 850, a memory 820, and a storage 830, each connected to an interconnect (bus) 840. The server computing system 110 may also include an I/O device interface 870 connecting I/O devices 875 (for example, keyboard, display, mouse, and/or other input/output devices) to the computing system 110. Further, in context of this disclosure, the computing elements shown in server computing system 110 may correspond to a physical computing system (for example, a system in a data center, a computer server, a desktop computer, a laptop computer, and/or the like) and/or may be a virtual computing instance executing within a hosted computing environment.

The CPU 860 may retrieve and execute programming instructions stored in memory 820, as well as store and retrieve application data residing in memory 820. The bus 840 may be used to transmit programming instructions and application data between the CPU 860, I/O device interface 870, storage 830, network interface 850, and memory 820. Note that the CPU 860 is included to be representative of, for example, a single CPU, multiple CPUs, a single CPU having multiple processing cores, a CPU with an associate memory management unit, and the like.

The memory 820 is included to be representative of, for example, a random access memory (RAM), cache and/or other dynamic storage devices for storing information and instructions to be executed by CPU 860. Memory 820 also may be used for storing temporary variables or other intermediate information during execution of instructions to be executed by CPU 860. Such instructions, when stored in storage media accessible to CPU 860, render server comput-

ing system 110 into a special-purpose machine that is customized to perform the operations specified in the instructions.

The storage 830 may be a disk drive storage device, a read only memory (ROM), or other static, non-transitory, and/or computer-readable storage device or medium coupled to bus 840 for storing static information and instructions for CPU 860. Although shown as a single unit, the storage 830 may be a combination of fixed and/or removable storage devices, such as fixed disc drives, removable memory cards, and/or optical storage, network attached storage (NAS), and/or a storage area-network (SAN).

Programming instructions, such as the cluster engine 120, the workflow engine 125, and/or the user interface engine 126, may be stored in the memory 820 and/or storage 830 in various software modules. The modules may be stored in a mass storage device (such as storage 830) as executable software codes that are executed by the server computing system 110. These and other modules may include, by way of example, components, such as software components, object-oriented software components, class components and task components, processes, functions, attributes, procedures, subroutines, segments of program code, drivers, firmware, microcode, circuitry, data, databases, data structures, tables, arrays, and variables.

Illustratively, according to an embodiment, the memory 820 stores a seed list 210, a cluster engine 120, a cluster list 250, a workflow engine 125, and a user interface engine 126 (as described with reference to the various figures above). The cluster engine 120 may include a cluster strategy 232-2. The particular cluster strategy 232-2 may include data bindings 237-1, 237-2, and 237-3, with which the cluster engine 120 may access the cluster data source 160. The workflow engine 125 may include a scoring strategy 442-1.

Illustratively, according to an embodiment, the storage 830 includes a cluster strategy store 230, data bindings store 835, a scoring strategy store 440, and one or more cluster analysis rules or criteria 880. As described above, the cluster strategy store 230 may include a collection of different cluster strategies 232, such as cluster strategy 232-2. For example, the cluster strategy store 230 may be a directory that includes the cluster strategies 232-1, 232-2 . . . 232-N as distinct modules. The scoring strategy store 440 may include a collection of different scoring strategies 442, such as scoring strategy 442-2, and may also be a directory of distinct modules. The data binding store 835 may include data bindings 237-1, 237-2 . . . 237-M, which may also be stored as distinct modules within a directory.

Although shown in memory 820, the seed list 210, cluster engine 120, cluster list 250, workflow engine 125, and the user interface engine 126, may be stored in memory 820, storage 830, and/or split between memory 820 and storage 830. Likewise, copies of the cluster strategy 232-2, data binding 237-1, 237-2, and 237-3, and scoring strategy 442-2 may be stored in memory 820, storage 830, and/or split between memory 820 and storage 830.

The network 150 may be any wired network, wireless network, or combination thereof. In addition, the network 150 may be a personal area network, local area network, wide area network, cable network, satellite network, cellular telephone network, or combination thereof. Protocols and components for communicating via the Internet or any of the other aforementioned types of communication networks are well known to those skilled in the art of computer communications and thus, need not be described in more detail herein.

As described above in reference to FIG. 1, the server computing system 110 may be in communication with one or

more data sources **160**. Communication between the server computing system **110** and the data sources **160** may be via the network **150** and/or direct. In an embodiment, an optional data aggregator/formatter device and/or system may aggregate various data from multiple data sources and/or may format the data such that it may be received by the server computing system **110** in a standardized and/or readable format. For example, when multiple data sources contain and/or provide data in various formats, the data aggregator/formatter may convert all the data into a similar format. Accordingly, in an embodiment the system may receive and/or access data from, or via, a device or system such as the data aggregator/formatter.

As described above, in various embodiments the system may be accessible by an analyst (and/or other operator or user) through a web-based viewer, such as a web browser. In this embodiment, the user interface may be generated by the server computing system **110** and transmitted to the web browser of the analyst. Alternatively, data necessary for generating the user interface may be provided by the server computing system **110** to the browser, where the user interface may be generated. The analyst/user may then interact with the user interface through the web-browser. In an embodiment, the user interface of the data analysis system may be accessible through a dedicated software application. In an embodiment, the client computing device **130** may be a mobile computing device, and the user interface of the data analysis system may be accessible through such a mobile computing device (for example, a smartphone and/or tablet). In this embodiment, the server computing system **110** may generate and transmit a user interface to the mobile computing device. Alternatively, the mobile computing device may include modules for generating the user interface, and the server computing system **110** may provide user interaction data to the mobile computing device. In an embodiment, the server computing system **110** comprises a mobile computing device. Additionally, in various embodiments any of the components and/or functionality described above with reference to the server computing system **110** (including, for example, memory, storage, CPU, network interface, I/O device interface, and the like), and/or similar or corresponding components and/or functionality, may be included in the client computing device **130**.

According to various embodiments, the data analysis system and other methods and techniques described herein are implemented by one or more special-purpose computing devices. The special-purpose computing devices may be hard-wired to perform the techniques, or may include digital electronic devices such as one or more application-specific integrated circuits (ASICs) or field programmable gate arrays (FPGAs) that are persistently programmed to perform the techniques, or may include one or more general purpose hardware processors programmed to perform the techniques pursuant to program instructions in firmware, memory, other storage, or a combination. Such special-purpose computing devices may also combine custom hard-wired logic, ASICs, or FPGAs with custom programming to accomplish the techniques. The special-purpose computing devices may be desktop computer systems, server computer systems, portable computer systems, handheld devices, networking devices or any other device or combination of devices that incorporate hard-wired and/or program logic to implement the techniques.

Computing devices of the data analysis system may generally be controlled and/or coordinated by operating system software, such as iOS, Android, Chrome OS, Windows XP, Windows Vista, Windows 7, Windows 8, Windows Server,

Windows CE, Unix, Linux, SunOS, Solaris, iOS, Blackberry OS, VxWorks, or other compatible operating systems. In other embodiments, the computing devices may be controlled by a proprietary operating system. Conventional operating systems control and schedule computer processes for execution, perform memory management, provide file system, networking, I/O services, and provide a user interface functionality, such as a graphical user interface (“GUI”), among other things.

In general, the word “module,” as used herein, refers to a collection of software instructions, possibly having entry and exit points, written in a programming language, such as, for example, Java, Lua, C or C++. A software module may be compiled and linked into an executable program, installed in a dynamic link library, or may be written in an interpreted programming language such as, for example, BASIC, Perl, or Python. It will be appreciated that software modules may be callable from other modules or from themselves, and/or may be invoked in response to detected events or interrupts. Software modules configured for execution on computing devices may be provided on a computer readable medium, such as a compact disc, digital video disc, flash drive, magnetic disc, or any other tangible medium, or as a digital download (and may be originally stored in a compressed or installable format that requires installation, decompression or decryption prior to execution). Such software code may be stored, partially or fully, on a memory device of the executing computing device, for execution by the computing device. Software instructions may be embedded in firmware, such as an EPROM. It will be further appreciated that hardware devices (such as processors and CPUs) may be comprised of connected logic units, such as gates and flip-flops, and/or may be comprised of programmable units, such as programmable gate arrays or processors. The modules or computing device functionality described herein are preferably implemented as software modules, but may be represented in hardware devices. Generally, the modules described herein refer to software modules that may be combined with other modules or divided into sub-modules despite their physical organization or storage.

Server computing system **110** may implement various of the techniques and methods described herein using customized hard-wired logic, one or more ASICs or FPGAs, firmware and/or program logic which, in combination with various software modules, causes the server computing system **110** to be a special-purpose machine. According to one embodiment, the techniques herein are performed by server computing system **110** in response to CPU **860** executing one or more sequences of one or more modules and/or instructions contained in memory **820**. Such instructions may be read into memory **820** from another storage medium, such as storage **830**. Execution of the sequences of instructions contained in memory **820** may cause CPU **840** to perform the processes and methods described herein. In alternative embodiments, hard-wired circuitry may be used in place of or in combination with software instructions.

The term “non-transitory media,” and similar terms, as used herein refers to any media that store data and/or instructions that cause a machine to operate in a specific fashion. Such non-transitory media may comprise non-volatile media and/or volatile media. Non-volatile media includes, for example, optical or magnetic disks, such as storage **830**. Volatile media includes dynamic memory, such as memory **820**. Common forms of non-transitory media include, for example, a floppy disk, a flexible disk, hard disk, solid state drive, magnetic tape, or any other magnetic data storage medium, a CD-ROM, any other optical data storage medium, any physical medium with patterns of holes, a RAM, a

PROM, and EPROM, a FLASH-EPROM, NVRAM, any other memory chip or cartridge, and networked versions of the same.

Non-transitory media is distinct from but may be used in conjunction with transmission media. Transmission media participates in transferring information between non-transitory media. For example, transmission media includes coaxial cables, copper wire and fiber optics, including the wires that comprise bus 840. Transmission media may also take the form of acoustic or light waves, such as those generated during radio-wave and infra-red data communications.

Various forms of media may be involved in carrying one or more sequences of one or more instructions to CPU 860 for execution. For example, the instructions may initially be carried on a magnetic disk or solid state drive of a remote computer. The remote computer may load the instructions and/or modules into its dynamic memory and send the instructions over a telephone or cable line using a modem. A modem local to server computing system 820 may receive the data on the telephone/cable line and use a converter device including the appropriate circuitry to place the data on bus 840. Bus 840 carries the data to memory 820, from which CPU 860 retrieves and executes the instructions. The instructions received by memory 820 may optionally be stored on storage 830 either before or after execution by CPU 860.

VII. Additional Example Applications

While financial fraud using credit card accounts is used as a primary reference example in the discussion above, the techniques described herein may be adapted for use with a variety of data sets and in various applications. Such applications may include, for example, financial fraud detection, tax fraud detection, beaconing malware detection, malware user-agent detection, other types of malware detection, activity trend detection, health insurance fraud detection, financial account fraud detection, detection of activity by networks of individuals, criminal activity detection, network intrusion detection, detection of phishing efforts, money laundering detection, and/or financial malfeasance detection. For example, information from data logs of online systems may be evaluated as seeds to improve cyber security. In such a case, a seed may be a suspicious IP address, a compromised user account, and the like. From the seeds, log data, DHCP logs, IP blacklists, packet captures, webapp logs, and other server and database logs may be used to create clusters of activity related to the suspicions seeds. Other examples include data quality analysis used to cluster transactions processed through a computer system (whether financial or otherwise). A number of examples of such applications are described in detail below in reference the various figures.

VIII. Example Generalized Method of the Data Analysis System

FIG. 9 is a flowchart of an example generalized method of the data analysis system, according to an embodiment of the present disclosure. In various embodiments, fewer blocks or additional blocks may be included in the process of FIG. 9, or various blocks may be performed in an order different from that shown in the figure. Further, one or more blocks in the figure may be performed by various components of the data analysis system, for example, server computing system 110 (described above in reference to FIG. 8).

As described above, and as shown in the embodiment of FIG. 9, the data analysis system may generate a seed or multiple seeds (block 910), may generate clusters based on

those seed(s) (block 920), may generate a score or multiple scores for each generated cluster (block 930), may generate a metascore for each generated cluster (block 940), and may optionally rank the generated clusters based on the generated metascores (block 950). In various embodiments, the data analysis system may or may not generate multiple scores for each cluster, may or may not generate metascores for each cluster, and/or may or may not rank the clusters. In an embodiment, the system may rank clusters based on one or more scores that are not metascores.

Further, as described above, the seeds may include one or multiple data items, and may be generated based on seed generation strategies and/or rules. Similarly, the clusters may include one or multiple data items related to a seed, including the seed, and may be generated based on cluster generation strategies and/or rules (including data bindings and/or searching and filtering are performed through, for example, a generic interface to various data sources). Scores and metascores may be determined based on attributes, characteristics, and/or properties associated with data items that make up a given cluster.

Example applications of the data analysis system, including methods and systems for identifying data items, generating data clusters, and analyzing/scoring clusters, are disclosed in the various related applications listed above and previously incorporated by reference herein.

IX. Cluster Analysis and Example Analysis User Interfaces

FIGS. 10A-10C and 11-22, described below, illustrate methods and user interfaces of the data analysis system, according to various embodiments, in which data clusters are automatically generated, analyzed, and presented to an analyst such that the analyst may quickly and efficiently evaluate the clusters. In particular, as described below the data analysis system may apply one or more analysis criteria or rules to the data clusters so as to generate human-readable "conclusions" (as described above, also referred to herein as "summaries"). The conclusions may be displayed in an analysis user interface through which the analyst may evaluate the clusters and/or access more detailed data related to the cluster. In an embodiment, a cluster type may be associated with each cluster, and may be determined according to the cluster strategy that generated the cluster. Further, the system may generate "alert scores" for the clusters which may be used to prioritize clusters displayed to the analyst.

The various methods and user interfaces described below in reference to FIGS. 10A-10C and 11-22 may be implemented by various aspects of the data analysis system (for example, the server computing system 110 and/or another suitable computing system) as described above. For example, clustering may be accomplished according to seed generation and clustering strategies and rules as implemented by, for example, the cluster/rules engine 120; cluster analysis may be accomplished according to analysis rules/criteria 880 as implemented by, for example, the cluster/rules engine 120; cluster scoring (for example, generation of alert scores) may be accomplished according to scoring strategies as implemented by, for example, the cluster/rules engine 120; and user interface may be generated and/or presented to the analyst by, for example, the user interface engine 126; among other aspects.

Additionally, in the methods described in reference to the flowcharts of FIGS. 10A-10B and 21 below, in various embodiments, fewer blocks or additional blocks may be included in the example methods depicted, or various blocks

45

may be performed in an order different from that shown in the figures. Further, in various embodiments, one or more blocks in the figures may be performed by various components of the data analysis system, for example, server computing system 110 (described above in reference to FIG. 8) and/or another suitable computing system.

a. Example Method of Cluster Analysis

FIG. 10A is a flowchart for an example method of data cluster analysis, according to an embodiment of the present disclosure. In FIG. 10A, blocks 910 and 920 of the flowchart proceed generally as described in reference to the flowchart of FIG. 9. For example, at block 910 seeds are generated according to one or more seed generation strategies. Examples of seed generation strategies are described in the various related applications listed above and previously incorporated by reference herein. Examples include identifying tax returns that are potentially fraudulent, identifying communications that are potentially associated with beaconing malware, and/or identifying emails potentially associated with phishing campaigns, among others. Further, at block 920 clusters are generated based on the one or more generated seeds and according to the one or more cluster generation strategies. Examples of cluster generation strategies (as mentioned above, also referred to herein as “cluster strategies,” “clustering strategies,” and/or “cluster generation rules”) are described in the various related applications listed above and previously incorporated by reference herein. Examples include strategies for financial fraud detection, tax fraud detection, beaconing malware detection, malware user-agent detection, other types of malware detection, activity trend detection, health insurance fraud detection, financial account fraud detection, detection of activity by networks of individuals, criminal activity detection, network intrusion detection, detection of phishing efforts, money laundering detection, and/or financial malfeasance detection, among others.

A cluster of data items generated according to a given clustering strategy (and its associated seed generation strategy or strategies) may be understood as having a “cluster type” (also referred to as a “data cluster type”) corresponding to that clustering strategy. For example, a particular clustering strategy may be referred to as “Tax Fraud,” because the clustering strategy relates to identifying clusters of data items related to potential tax fraud. A cluster of data items generated according to that clustering strategy may therefore have a “cluster type” of “Tax Fraud.” In another example, a cluster generated by an “Internal Phishing” clustering strategy (and its associated seed generation strategy or strategies) has a cluster type of “Internal Phishing.”

At block 1002 of the system generates “alerts” for each of the clusters. An “alert” includes various types of information related to the cluster that may be useful to an analyst in evaluating the importance or criticality of the cluster in the context of a particular investigation. Generating an alert may include applying various cluster analysis rules or criteria to analyze the cluster and so as to generate human-readable cluster conclusions, as mentioned above. Generating an alert may further include generating an alert score for the cluster. Details regarding generation of alerts are described below in reference to FIG. 10B.

At block 1004 of FIG. 10A, a cluster analysis user interface is provided to the user (for example, an analyst). FIGS. 11-20 and 22, described below, include examples of cluster analysis user interfaces of the data analysis system. As described below, a user interface may include a listing of alerts, each alert corresponding to a particular generated and analyzed cluster. The alerts may be organized and grouped according to cluster types. Further, the analyst may view a user interface

46

including detailed information related to each alert, including the human-readable conclusions, the alert scores, and various detailed data related to the clusters. For example, in a given alert the analyst may be provided with a name of the cluster, a cluster strategy by which the cluster was generated (also referred to as the cluster type), a list of generated conclusions, and/or one or more lists and/or tables of data related to the cluster. The one or more lists and/or tables of data related to the cluster may be drawn from the data items of the cluster, and may be filtered by the analyst according to time and/or type of data.

At block 1006, the system regenerates previously generated clusters. In various implementations the data items from which seeds are selected/generated and from which clusters are generated may change after a cluster is generated. In the example of tax fraud detection, additional tax return data items may be received, or additional phone number data items may be received that relate to a person in a previously generated cluster. Such information may have been included in a cluster if it had been available at the time the cluster was created. Accordingly, the system may regenerate clusters so as to include the data items and/or other information that has become available since the last time the cluster was generated. After, or in response to, a cluster being regenerated, the system reanalyzes the cluster and may, in an embodiment, generate an alert for the regenerated and reanalyzed cluster (as indicated by the arrow back to block 1002). In another embodiment, as described below in reference FIG. 21, when a given cluster is regenerated, a previously generated alert for that cluster may be updated or, alternatively, a new alert may be generated including a link to the previously generated alert.

In an embodiment, as shown at block 1006a, clusters may be regenerated on a schedule. For example, the system may be configured to regenerate clusters after a particular number of seconds, minutes, hours, or days, or at particular times every hour or day. In another embodiment, as shown at block 1006b, clusters may be regenerated as needed, such as in response to the system detecting one or more changes in data items and automatically executing a cluster regeneration process. For example, the system may be configured to automatically regenerate clusters when it detects that new data items (and/or other information) are received by the system, new data items (and/or other information) related to a cluster (and/or potentially related to a cluster) are received by the system, new data items (and/or other information) connected to a cluster or a data item in a cluster is received by the system, an analyst logs into the system, and/or an analyst views a cluster. In another embodiment, as shown at block 1006c, clusters may be regenerated on demand. For example, clusters may be regenerated when requested by an analyst (via, for example, a user interface of the system).

In any of the embodiments of blocks 1006a, 1006b, and 1006c, all clusters may be regenerated or portions of clusters may be regenerated, in any combination. For example, clusters associated with a particular clustering strategy may be generated on a particular schedule, while clusters associated with a different clustering strategy may be generated on a different schedule (and/or as needed and/or on demand). In another example, individual clusters may be regenerated, or other relationships among clusters may be used to determine which clusters are to be regenerated at a given time.

At optional block 1008, clusters are merged as described above. For example, if a regenerated cluster includes a data item also included in a different cluster, the regenerated cluster and the different cluster may optionally be merged. In the embodiment of FIG. 10A, only clusters generated according

to the same clustering strategy (for example, having the same cluster type) may be merged. In this embodiment, alerts generated for clusters having different cluster types may be linked even if the clusters are not merged, as described below in reference to FIGS. 21 and 22. In alternative embodiments,

b. Example Method of Alert Generation

FIG. 10B is a flowchart of an example method of alert generation for a particular data cluster, according to an embodiment of the present disclosure. At block 1022, the system accesses data, including data items and related meta-data and other information, of the data cluster. As described below, this accessed cluster data is analyzed to generate the human-readable conclusions, the alert scores, and may be included and organized in the user interface of the alert. At block 1024, the system determines the cluster type of the data cluster. As mentioned above, a data cluster generated according to a given clustering strategy (and its associated seed generation strategy or strategies) may be understood as having a “cluster type” (also referred to as a “data cluster type”) corresponding to that clustering strategy.

At block 1026, having determined the cluster type of the data cluster, the system accesses one or more cluster analysis rules or criteria associated with that cluster type. As various data clusters may be generated according to different clustering strategies, and each of the clustering strategies may be associated with differing types of investigations, the analysis rules or criteria used to analyze the clusters vary according to the cluster types and their respective associated types of investigations.

At block 1028, the system analyzes the data cluster based on the accessed analysis rules/criteria. The cluster data is then evaluated by the system (for example, by the cluster/rules engine 120) according to the analysis rules/criteria. Many examples of cluster analysis according to various clustering strategies are described in the various related applications listed above and previously incorporated by reference herein. In the various examples, analysis of clusters may be described in the context of cluster scoring (for example, generating of clusters scores and/or metascores). For example, in U.S. patent application Ser. No. 14/139,628, cluster data is scored and/or analyzed in various contexts including, among others:

Tax Fraud Detection, in which clusters are analyzed to determine a number of known fraudulent returns in a cluster, a number of first-time filers in the cluster, and/or a mismatch between reported incomes in the cluster, among others.

Beaconing Malware Detection, in which clusters are analyzed to determine a number of known bad domains in a cluster, an average request size in the cluster, and/or a number of requests blocked by a proxy in the cluster, among others.

Additional examples are described below in reference to FIGS. 24A-24C, 25A-25B, 26A-26C, and 27A-27C, in which cluster data is scored and/or analyzed in various contexts including:

Internal and External Phishing, in which clusters are analyzed to determine a most common email subject of emails in the cluster, numbers of emails in the cluster sent within particular time periods, and/or number of recipients of emails in the cluster, among others.

Internal and External Threat Intel, in which clusters are analyzed to determine a number of URLs in the cluster referenced by an analyzed malware data item, a percentage of traffic in the cluster categorized as likely malicious, and/or a highest organizationally hierarchical

position of a person in the cluster associated with a malicious connection, among others.

IDS (Intrusion Detection System), in which clusters are analyzed to determine a time spanned by alert notices in the cluster and/or a number of alert notices associated with particular IP addresses, among others.

i. “Conclusions”

At block 1030 of FIG. 10B, the system generates one or more conclusions for the analyzed data cluster based on the cluster analysis. As described above, the generated conclusions (also referred to herein as summaries) comprise compact, human-readable phrases or sentences that provide highly relevant, and easily evaluated (by a human analyst), information regarding the data in the cluster (for example, data items and metadata). The conclusions may be useful to an analyst in evaluating the importance or criticality of the cluster in the context of a particular investigation. As with the analysis rules/criteria described above, each cluster type may be related to a set of conclusions appropriate to the type of investigation associated with the cluster type. FIG. 10C illustrates various example templates for conclusions (also referred to herein as “conclusion templates”) associated with various types of data clusters, according to an embodiment. For example, five cluster types (which are each associated with various seed generation, clustering, and scoring strategies) are included in the example embodiment of FIG. 10C: Internal Phishing, External Phishing, Internal Threat Intel, External Threat Intel, and IDS (short for Intrusion Detection System). Each of the example cluster types is associated with one or more conclusion templates, as shown in the right column of the table of FIG. 10C. The conclusion templates include fields (indicated by the symbols < and >) into which cluster information, obtained as a result of the cluster analysis, is inserted when the conclusion is generated.

For example, in reference to the embodiment of FIG. 10C, for the cluster type “External Phishing,” a conclusion template is “This campaign consists of <m> emails submitted to external Abuse,” where <m> indicates a field to be filled in by the system based on the cluster analysis, and “external Abuse” may refer to an email address or box. In generating this conclusion, the system accesses the relevant set of conclusions (for example, conclusions associated with the type of the cluster analyzed) and inserts relevant cluster analysis data into each of the conclusions (for example, “This campaign consists of 25 emails submitted to external Abuse”). In another example, for the cluster type “External Threat Intel,” a conclusion template is “<k>% of proxy traffic was blocked, and <l>% was marked as malicious by Proxy,” where <k> and <l> indicate fields to be filled in by the system based on the cluster analysis. In generating this conclusion, the system accesses the relevant set of conclusions (for example, conclusions associated with the type of the cluster analyzed) and inserts relevant cluster analysis data into each of the conclusion templates (for example, “10% of proxy traffic was blocked, and 7% was marked as malicious by Proxy”).

In an embodiment, conclusion templates, such as those listed in the table of FIG. 10C, may be manually generated by humans based on a determination of information likely to be helpful to an analyst in evaluating alerts/clusters. The manually generated conclusion templates associated with respective cluster types may then be automatically accessed by the system (e.g., after automatically determining which conclusion templates are applicable), relevant data may be inserted into any indicated fields, and conclusions may then be automatically generated based on the selected conclusion template(s) and presented on a user interface (as described below). In another embodiment, the system may automati-

cally use heuristics to generate conclusion templates that may then be presented by the system. In this example, the system may determine, over time, information most useful to analysts, and thereby generate conclusion templates and conclusions based on that useful information.

In an embodiment, a predefined group of conclusions may be associated with each cluster type. In this embodiment, all conclusions in the relevant group may be generated and presented in the user interface for each respective alert. In another embodiment, various conclusions may be associated with each cluster type, and the system may determine particular conclusions, based on the cluster analysis, to generate and present in the user interface. In this embodiment, the system may select particular conclusions based on a likelihood that the particular conclusions will be helpful to the analyst in evaluating the cluster. For example, when a cluster does not have any data items (and/or other information) of a particular type that are enumerated (and/or otherwise evaluated) in a particular conclusion, that particular conclusion may not be displayed to the analyst. Alternatively, the system may indicate to the analyst that the particular conclusion is not applicable to the cluster.

In an embodiment, conclusions may be unique to each cluster type. In another embodiment, conclusions may be applicable to multiple cluster types.

In an embodiment, a conclusion may not express an opinion, but may only provide factual information. For example, "Less than 1 MB of data was exchanged with the following URL: <http://example.com>." In another embodiment, a conclusion may express an opinion if a judgment threshold is provided (for example, some factual basis for the opinion), but not otherwise. For example, an appropriate conclusion may be "Only a small amount of data, 0.7 MB, was exchanged with the following URL: <http://example.com>," while an inappropriate conclusion may be "Only a small amount of data was exchanged with the following URL: <http://example.com>." In various embodiments, conclusions generated by the system provide factual and/or opinion information to the analyst in the context of a particular investigation and/or cluster/alert type.

In an embodiment, each conclusion is limited to a particular number of words, for example, 10, 15, or some other number. In an embodiment, each user interface associated with an alert (as described below) displays between one and some other number of conclusions, for example, 2, 3, 4, 5, among others.

ii. "Alert Score"

Turning again to the embodiment shown in FIG. 10B, at block 1032, the system generates an alert score for the analyzed data cluster based on the cluster analysis. As described above, the alert score may be the same as, similar to, and/or based on any of the scores, metascoring, and/or conclusions described herein. An alert score may provide an initial indication to an analyst of a likelihood that a cluster/alert is important or critical in the context of a particular investigation (for example, a degree of correlation between characteristics of the cluster/alert and the analysis rules/criteria). As described below, the alert score is represented in the analysis user interface by an indicator, icon, color, and/or the like. An analyst may sort alerts/clusters based on the alert scores so as to enable an efficient investigation of more important alerts/clusters first.

In an embodiment, the alert score may be a metascoring, and may be one of multiple values. For example, the alert score may be one of three values corresponding to, for example, a high alert, a medium alert, or a low alert. In other embodiments, the alert score may be partitioned into more or fewer

values. Examples of various scores and metascoring associated with various cluster strategies are described in the various related applications listed above and previously incorporated by reference herein. For example, in U.S. patent application Ser. No. 14/139,628, example cluster metascoring are described in the contexts of tax fraud detection, beaconing malware detection, malware user-agent detection, and activity trend detection.

As mentioned above, in an embodiment, the alert score may be binned into one of three bins corresponding to a high alert, a medium alert, or a low alert. Each alert level may be associated with an indicator, icon, color, and/or the like. For example, a high alert may be associated with red (and/or another color), a medium alert may be associated with orange (and/or another color), and a low alert may be associated grey (and/or another color).

In an embodiment, the cluster alert score is determined based on and conveys both a determined importance/criticality (for example, a metascoring comprising scores showing a high number of data items may indicate likely fraud) and a confidence level in the determined importance/criticality. For example:

A high alert may be indicated when:

an importance metascoring is above a particular threshold (for example, greater than 60%, or some other percent or number), AND a confidence level is above a particular threshold (for example, greater than 70%, or some other percent or number).

A medium alert may be indicated when:

an importance metascoring is below a particular threshold (for example, less than 60%, or some other percent or number), AND a confidence level is above a particular threshold (for example, greater than 70%, or some other percent or number), OR

an importance metascoring is above a particular threshold (for example, greater than 60%, or some other percent or number), AND a confidence level is below a particular threshold (for example, less than 30%, or some other percent or number).

A low alert may be indicated when:

either an importance metascoring is below a particular threshold (for example, less than 60%, or some other percent or number), OR a confidence level is below a particular threshold (for example, less than 30%, or some other percent or number).

In other embodiments, other criteria may be used to determine alert levels to provide to the end user, possibly based on additional or fewer parameters than discussed above. In some examples, alerts are associated with ranges of importance metascoring and/or confidence levels, rather than only a minimum or maximum level of particular scores as in the examples above.

In an embodiment, a confidence level may be determined based on a false positive rate. The false positive rate may be based on, for example, historical information indicating how frequently other clusters having similar fraud indicators (for example, indicators used in the determination of the importance metascoring) have been determined, after human analysis, to be critical or not consistent with the importance metascoring. The false positive rate may also (or alternatively) be based on, for example, information provided from third-parties, such as blacklists that include a likelihood that any item on the blacklist is a false positive.

As mentioned above, in an embodiment the alert score may be based on one or more cluster scores and/or the analysis rules/criteria. In this embodiment, a high alert score may indicate a high degree of correlation between characteristics

51

(for example, data and metadata) of the cluster and the analysis rules/criteria (that may, for example, indicate a likelihood of fraud, among other indications). Similarly, a low alert score may indicate a high degree of correlation between characteristics of the cluster and the analysis rules/criteria.

c. Example Analysis User Interfaces

FIGS. 11-20 illustrate example data cluster analysis user interfaces of the data analysis system, according to embodiments of the present disclosure. In various embodiments, aspects of the user interfaces may be rearranged from what is shown and described below, and/or particular aspects may or may not be included. However, the embodiments described below in reference to FIGS. 11-20 provides example analysis user interfaces of the system.

FIG. 11 illustrates a user interface 1102 of the system in which various indicators of alerts associated with various types of clusters are displayed, according to an embodiment. The user interface 1102 includes a panel 1104 including a listing of various cluster types 1106 (which are each associated with respective clustering strategies). Selection of one of the cluster types 1106 results in a display of indications of associated alerts in the panel 1112. In FIG. 11, selection of "All" 1108 causes display of a combined list of indicators associated with all types of clusters in the panel 1112. Indicator 1110 shows a number of alerts among all the cluster types. In the panel 1112, at 1114 it is indicated that the present view is the "Inbox." The Inbox includes indications of alerts that have not yet been "Archived" by the analyst (as described below). Alternatively, the Inbox may show indications of alerts that have not yet been viewed by the analyst.

At 1116 an indication of an alert is shown. As discussed above, each listed alert corresponds to a particular data item cluster that has been generated, analyzed, and scored. Various details related to the alert are displayed including an alert title (for example, "!! Activity summary for Acct#1074911"), an indication of a time 1118 when the event associated with the alert occurred (for example, "1 hour ago"), and an indication of the cluster type 1120 (for example, "SYNTHETICS"). The alert title may be a single, human-readable summary phrase or sentence, and may be generated similar to the generation of conclusions described above, and/or may be (or include) one of the conclusions described above. In the example shown, the alert 1116 is related to identification of fraudulent bank accounts, and the alert title indicates the number of the primary bank account associated with the cluster. Additionally, the "!!" symbol shown at the beginning of the alert title provides an indication of the alert score of the alert. In the example shown, a "!!" indicated a medium risk level, a "!!!" indicates a high risk level, and no symbol indicates a low risk level. In other embodiments the alert level of an alert may be indicated by an icon and/or coloring of the alert indicator, among other indications. The analyst may select any of the listed alert indicators to view additional detail related to the selected alert. In an embodiment, the list of alert indicators may automatically be sorted according to one or more criteria, for example, the alert score. In an embodiment, the analysis may choose to sort the list of alert indicators as desired. In an embodiment, the time 1118 may be a time when the alert was generated, rather than the time the event associated with the alert occurred. In another embodiment, the time 1118 may include both the time the alert was generated and the time the event associated with the alert occurred.

FIG. 12 illustrates a user interface 1202 of the system in which a particular selected alert is displayed, according to an embodiment. The upper portion 1203 of the user interface may be colored to correspond to the alert score, as described above. At 1204 the alert title is displayed. A unique icon 1205

52

associated with the cluster types may be displayed. At 1206, an indication of the cluster type is given, as well as a unique identifier of the alert (for example, "#116," which may be useful for further investigation, note taking, and/or sharing by the analyst). At 1208 various conclusions (generated as described above) associated with the cluster are displayed. For example, in the cluster represented by the alert shown, the conclusions indicate that there have been no money transfers to other accounts, there are 13 transactions, the largest transaction is \$9,897.61, and 2 online accounts have been accessed by 29 computers. Such information may be helpful to an analyst in evaluating whether or not the alert includes accounts associated with fraudulent identities (also referred to as synthetic identities).

Selectable buttons 1210 and 1212 (and/or other user interface elements) are displayed by which the analyst may access detailed cluster data. For example, the analyst may select "Latest Online Account Logins" button 1210 to view a listing of most recent account login data panel 1214. Similarly, the analyst may select "Latest Transactions" 1212 to view a listing of transaction data in the panel 1214. Additional buttons or controls may be included in the display such that the analyst may view other data related to the cluster. As shown, the data displayed in the panel 1214 may be organized in a table including columns and rows. Data displayed may be drawn from various data items and/or other information included in the cluster. The particular buttons (such as buttons 1210 and 1212) displayed in the alert may be defined by the clustering strategy and/or another set of rules related to the cluster type. FIG. 13 shows the same alert as shown in FIG. 12, however the latest transactions button 1212 has been selected by the analyst, such that the information in panel 1214 is updated to show a listing of most recent transactions. In an embodiment, information shown in the panel 1214 may be automatically sorted chronologically from most recent event. Further the analyst may select the button 1302 to view further additional cluster data.

In other embodiments, the user interface may include links (for example, via buttons or other user interface elements) to relevant cluster information internal to an organization using the data analysis system, external to the organization, and/or other types information.

FIG. 14 illustrates a user interface 1402 of the system that is displayed when the show logs button 1302 (of FIG. 13) is selected, according to an embodiment. The user interface includes various cluster data and information 1404 organized in a table, a dropdown list of data types or sources 1406, a time filter 1408, and the button 1302 that may be selected to go back to the alert display of FIG. 12 or 13. The cluster data and information 1404 may be drawn from various data items and/or other information included in the cluster. The table shown is a stacked table, meaning that multiple differing types of data are displayed in the table, and the types of data displayed in a given column or row of the table may change within the given column or row. For example, as the dropdown 1406 indicates that all data types are displayed, the top portion of the table, as indicated by the left-most column, includes Address data items, the next portion of the table (below the top portion) includes Transaction data items, the next portion of the table includes Account data items, the next portion of the table includes Online Account data items, and the bottom portion of the table includes Customer data items. The analyst or other user may scroll down the table to view additional table entries, and/or may scroll horizontally to view additional columns of the table. In various embodiments the table may or may not be sorted by default in a chronological order, and the columns may or may not be arranged such

53

that the first column for each data type is a timestamp. In an embodiment, information displayed in the table is raw data drawn from entries associated with data items of the cluster.

FIG. 15 illustrates the same user interface as shown in FIG. 14, but shows various changes made by the analyst, according to an embodiment. For example, the analyst has selected the dropdown box 1406 to view the various types of data that may be selected. Further, the analyst has moved a starting-time indicator 1502 on the time filter 1408. Moving the starting-time indicator 1502 causes the data displayed in the table to be filtered to include only data that was produced and/or relates to items or events that occurred within a time span indicated by the starting-time indicator 1502 and an ending-time indicator 1504.

FIG. 16 illustrates the same user interface as shown in FIG. 14, but shows various changes made by the analyst. For example, the analyst has selected to view only Transaction data items via the dropdown box 1406. Further, the analyst has adjusted the time filter 1408 to filter that data items for a different particular time span. Accordingly, the table 1602 only displayed Transaction information related to the specified time span.

FIG. 17 illustrates the same user interface as shown in FIG. 14, but shows that the analyst may further filter the data displayed in the table by values in any of the columns. For example, a Type dropdown menu 1702 may be used by the analyst to specify particular types of transactions that are to be displayed in the table, such that other types of transactions are not displayed. The analyst may specify multiple types by selection and/or text input, and may selectively remove types that are selected.

FIG. 18 illustrates a user interface 1802 similar to the user interface of FIG. 11, according to an embodiment. In the user interface of FIG. 18, at 1804 the user has selected to view only indications of alerts of the type "Synthetics." Additionally, FIG. 18 illustrates that the analyst may select multiple indications of alerts, as shown at 1806, such that multiple alerts may be "archived" simultaneously by selection of the archive button 1808. Archiving alerts causes the alerts to be removed from the "Inbox" display. As shown in user interface 1902 of FIG. 19, the analyst may select to view "archived" alerts via the dropdown box 1904. Archived alerts are displayed in a list similar to the list of alerts provided in the Inbox. In an embodiment, archiving of alerts enables an analyst to indicate that they have reviewed a particular alert. The analyst may move the alert from the archive back to the inbox. Further, in other embodiments, alerts may be moved to additional categories (default and/or user defined), for example, a "Starred" category may be available. Archived alerts may automatically be moved back into the inbox when new data items are added to a cluster associated with an archived alert, such as when the cluster is regenerated, for example.

FIG. 20 illustrates a user interface 2002, similar to the user interface of FIG. 11, in which the left panel 1104 has been collapsed (as indicated by 2004) to provide a more streamlined display for the analyst, according to an embodiment.

In an embodiment, the alert user interface, for example the user interface of FIG. 12, may include user interface elements (such as buttons) selectable by the analyst to cause the system to archive an alert, categorize an alert, change an alert level, and/or share an alert with other analysts. In an embodiment, the alert user interface may include a button to add the cluster data items of a graph, as described in various related applications listed above and previously incorporated by reference herein. Further, the system may enable an analyst viewing a graph of data items to go to alerts representing clusters in which that data item is included.

54

In an embodiment, the analysis user interface, for example the user interface of FIG. 11, may include further details related to each of the indicated alerts. For example, the user interface of FIG. 11 may include conclusions associated with each of the listed alert indications. Providing data to the analyst in this way may enable the analyst to efficiently evaluate clusters without necessarily viewing the alert user interface.

In an embodiment, the analysis user interface, for example the user interface of FIG. 11, may include, in the list of indications of alerts, indications of events of interest to the analyst but generated by other processed. For example, the list may include indications of notices generated by third-party software (for example, a virus scanner).

d. Linking of Related Alerts/Clusters

FIG. 21 is a flowchart of an example method of linking related alerts or data clusters, according to an embodiment of the present disclosure. As described above, when clusters are regenerated, if two clusters of the same type have common data items, the two cluster of the same type may then be merged. However, when two clusters having different cluster types include common data items, they are not generally merged. In order to notify the analyst that two data clusters of different types have common data items, the example method of FIG. 21 may be executed by the system. Such a notification may advantageously enable an analyst, for example, to find additional connections in the context of an investigation. For example, the analyst may discover that an item of malware associated with a malware cluster is hosted at a website that is linked to by phishing emails in a phishing cluster.

In the example method of FIG. 21, at block 2102, the system finds or determines clusters of different cluster types (for example, that were generated according to different clustering strategies) that have common data items (and/or other information). At optional block 2104, a link between the related clusters/alerts may be generated. FIG. 22 illustrates an example data cluster analysis user interface 2202 in which related alerts or data clusters are linked to one another, according to an embodiment of the present disclosure. As shown, at 2404 links from the current alert/cluster to two other related alerts/clusters is provided. The analyst may then select one of the links (for example, either "Cluster ABC" or "Cluster XYZ") to view the alert pertaining to that cluster. In an embodiment, an indication of the common data items among the clusters is provided in the user interface.

Turning again to the example method of FIG. 21, at optional block 2106 the clusters/alerts may be merged. For example, rather than simply linking among related alerts (as in FIG. 22), the system may combine the alerts into a single alert user interface.

Further, in the example method of FIG. 21, at optional block 2108, the analyst may be notified when two clusters/alerts are linked or related. For example, the analyst may be notified via a popup message displaying in the analysis user interface, via an email or other message, and/or via any other appropriate communications method.

e. Regenerated Clusters/Alerts

In an embodiment, when a cluster is regenerated, as described above with reference to block 1006 of FIG. 10A, an alert may be updated, the analyst may be notified, and/or a new alert may be generated. FIG. 23 is a flowchart of an example method of updating alerts in response to cluster regeneration, according to an embodiment of the present disclosure. At block 2302 of the example method shown, a cluster has been regenerated. At block 2304, the system determines whether any changes have been made to the cluster (for example, any new data items added to the cluster). If not, then

55

at block 2306 the alert corresponding to the cluster is not updated and the method ends. If so, then at block 2308 the system determines whether the analyst has viewed and/or archived the alert corresponding to the cluster. If not, then at block 2310 the alert is updated such that the cluster analysis is rerun, and the alert data (for example, the conclusions) is regenerated, on the new cluster including the new data items. In this block, as the analyst has not previously interacted with the alert, no notifications regarding changes to the alert/cluster are provided. If the analyst has viewed and/or archived the alert, then at blocks 2312 and 2314 the alert may be updated, changes to the alert may be shown in the alert user interface, and/or a new alert may be generated and links between the new and old alerts may be generated and provided in the alert user interfaces. For example, if the analyst was to select an old alert that had been superseded due to cluster regeneration, the system may automatically forward the analyst to the new alert and display a message such as "You have been redirected to the most recent version of this alert. Return to alert 277." Selection of "Return to alert 277" may cause the old alert to be displayed, where a message may be included such as "There is a more recent version of this alert," (which may link to the new alert).

In an embodiment, when regenerated clusters of a same cluster type are merged, alerts corresponding to those previous two clusters may be merged and updates may be displayed, and/or a new alert may be generated (and linked to from the old alerts) as described above.

In an embodiment, the system may provide a data feed including timely updates (including analysis information) on any changes to any previously generated clusters, and/or any newly generated clusters.

X. Example Applications of the Data Analysis System

FIGS. 24A-24C, 25A-25B, 26A-26C, and 27A-27C, described below, illustrate a number of example applications of the data analysis system, according to embodiments of the present disclosure, in which, for example, seeds are generated, clusters are generated, clusters are analyzed, alerts are generated, conclusions are generated, alert scores are generated, and/or analysis user interfaces are generated and/or presented. For example, FIGS. 24A-24C and 25A-25B illustrate embodiments of the data analysis system as applied to detection of malware threats, FIGS. 26A-26C illustrate embodiments of the data analysis system as applied to network intrusion activities, and FIGS. 27A-27C illustrate embodiments of the data analysis system as applied to phishing threat detection.

In general the methods described below in reference to FIGS. 24A-24B, 25A-25B, 26A-26B, and 27A-27B illustrate various example embodiments of aspects of the data analysis system as described above in reference to, for example FIG. 10A. For example, FIGS. 24A, 25A, 26A, and 27A-27B illustrate example methods of seed generation (according to various seed generation strategies) corresponding to block 910 of FIG. 10A; and FIGS. 24B, 25B, 26B, and 27A-27B illustrate example methods of cluster generation (according to various cluster generation strategies) corresponding to block 920 of FIG. 10A.

As described above in reference to the embodiments of FIGS. 10A-10B, after data clusters of one or more cluster types are generated, the clusters may be analyzed, alerts may be generated, and analysis user interfaces may be generated and presented to an analyst (for example, as shown in the example user interfaces of FIGS. 11 and 12). In the embodi-

56

ment of FIG. 10B described above, a given cluster is analyzed, conclusions are generated, and an alert score is generated according to analysis rules/criteria (and/or other rules/criteria, for example, scoring rules) specific to a type of the given cluster. FIG. 10C described above illustrates example conclusions that may be used in the example applications of the data analysis system described below in reference to FIGS. 24A-24C, 25A-25B, 26A-26C, and 27A-27C. Further, FIGS. 24C, 26C, and 27C illustrate additional embodiments of user interfaces showing alerts, similar to the user interface of FIG. 12 described above, as applied to the described example applications of the data analysis system.

In the example methods described below in reference to each of FIGS. 24A-24B, 25A-25B, 26A-26B, and 27A-27B, in various embodiments, fewer blocks or additional blocks may be included in the example methods depicted, or various blocks may be performed in an order different from that shown in the figures. Further, in various embodiments, one or more blocks in the figures may be performed by various components of the data analysis system, for example, server computing system 110 (as described above), the cluster/rules engine 120 (as described above), the user interface engine 126 (as described above), other components or aspects of the data analysis system, and/or another suitable computing system.

XI. Example Application of the Data Analysis System to Malware Threat Detection

FIGS. 24A-24C and 25A-25B illustrate various methods and user interfaces of the data analysis system as applied to malware threat detection, according to embodiments of the present disclosure. Malware may include any software program (and/or group of software programs) installed on a computer system and/or a network of computer systems maliciously and/or without authorization. When executed, an item of malware may take any number of undesirable actions including, for example, collection of private or sensitive information (for example, personal data and information, passwords and usernames, and the like), transmission of the collected information to another computing device, destruction or modification of data (for example, accessing, modifying, and/or deleting files), communication with other malware, transmission or replication of malware onto other connected computing devices or systems, transmission of data so as to attack another computing device or system (for example, a Distributed Denial of Service Attack), and/or hijacking of processing power, just to name a few. In most cases such malware infects a computing device via a network connection (for example, a connection to the Internet), and communicates with another computing device or system (for example, another Internet-connected computing device) to accomplish its purpose. Oftentimes malware is well hidden in the infected computing device such that it may not be detectable to an average user of the computing device.

Detection and removal of malware from infected computing devices and/or systems is a highly desirable, but oftentimes challenging task. Detection of malware is of particular importance to organizations (for example, businesses) that maintain internal networks of computing devices that may be connected to various external networks of computing devices (for example, the Internet) because infection of a single computing device of the internal network may quickly spread to other computing devices of the internal network and may result in significant data loss and/or financial consequences. Traditional anti-malware software applications may detect previously known malware on a single computing device, but a network of computing device may remain vulnerable and no

further information and/or related data items may be provided to, for example, a user of the computing device.

The data analysis system may be used, as described below, to automatically, efficiently, and effectively detect malware activities on a network, and identify and cluster various data items related to the detected malware. The clustered data items may include, for example, particular computing devices infected, users of those computing devices, and/or the like (and as described below). Such a data cluster may be analyzed by the data analysis system and an alert (including various conclusions) may be generated and presented to an analyst. The analyst may then efficiently review the alert, accurately determine whether the alert indicates a malware threat (and assess a level of the threat), and take action to remove the malware from all infected computing devices. Further, as the alert/cluster includes users (for example, data items representing users) associated with infected computing devices, the analyst may determine particular users and user behaviors that may have contributed to the malware infection, and mitigate future risks (through, for example, user education).

Assessment of the level of the threat associated with a particular Alert (for example, generation of an alert score and/or evaluation by an analyst) may be enabled by clustering of organizational and/or hierarchical positions or groups (also referred to herein as “bands”) associated with the clustered users. The organizational and/or hierarchical positions or groups, or “bands,” may be specific to a particular embodiment of the data analysis system and/or a particular organization for which data items are being analyzed and clustered. For example, “bands” for a particular organization may include “Senior Management” (including CEOs, COOs, and the like), “Middle Management” (including Vice Presidents and the like), “Lower Management” (including supervisors and the like), and “Staff” (including other employees of the organization). Other organizations may include other “bands,” may include more or fewer “bands,” and/or users may be organized into bands differently. Accordingly, in various embodiments of the data analysis system as applied to malware threat detection, various implementations of “bands,” or organizational and/or hierarchical positions or groups may be used, associated with users, and/or clustered. In some embodiments, particular organizational positions are associated with each user (for example, data item representing a user), and the data analysis system determines a band associated with the organizational position based on one or more rules or criteria.

In addition to the description below, examples of the data analysis system as applied to detection of malware (including seed generation, clustering, and cluster analysis and scoring) are described in U.S. patent application Ser. No. 14/139,603, previously incorporated by reference herein. Aspects of the examples of clustering, analysis, and scoring described in U.S. patent application Ser. No. 14/139,603 may be applied to the embodiments of the data analysis system as applied to malware threat detection described below.

As described below, the data analysis system may be used in a network environment in which an internal network is in communication with an external network. The system may be used to determine whether any computer systems of the internal network have been infected by malware that is communicating with computer systems of the external network. Various computerized devices may be included in the internal network that may be capable of capturing and/or logging data traffic between the internal network and the external network including, for example, network routers, proxy devices, and/or switches.

a. Malware Threat Detection: Seed Generation

FIGS. 24A and 25A are flowcharts showing example methods of seed generation of the data analysis system as applied to malware threat detection, according to embodiments of the present disclosure. The seed generation methods of FIGS. 24A and 25A may generally be understood to correspond to block 910 (generate seed(s)) of the generalized process of FIG. 10A. FIG. 24A relates to an embodiment of the data analysis system in which seeds are generated based on received suspected malware files, while FIG. 25A relates to an embodiment of the data analysis system in which seeds are generated based on external threat lists.

i. Internally Identified Seeds

As mentioned above, according to an embodiment seeds may be generated by the data analysis system in the context of malware threat analysis based on received suspected malware files. Such an embodiment is described below in reference to FIG. 24A.

Turning to the embodiment of FIG. 24A, at block 2402 one or more suspected malware files (also referred to herein as file data items) are received by the system. The file data items (suspected malware files) may be submitted to the system, for example, via a user interface and as described in reference to block 102 of FIG. 1 of U.S. patent application Ser. No. 14/473,860, titled “MALWARE DATA ITEM ANALYSIS,” previously incorporated by reference herein. In other embodiments, file data items may be submitted for analysis based on other factors, such as when a file data item is stored, accessed, and/or updated on a storage device of the system. At block 2404 (of FIG. 24A), a basic analysis of the file data item is initiated by the system. Examples of the basic analysis initiated by the system are described in reference to block 106 of FIG. 1 of U.S. patent application Ser. No. 14/473,860, titled “MALWARE DATA ITEM ANALYSIS,” previously incorporated by reference herein. For example, MD5 and SHA-1 hashes of the file data item may be determined, and/or a file size of the file data item may be determined, among other analyses. At block 2406 (of FIG. 24A), an external analysis of the file data item is initiated by the system. Examples of the external analysis initiated by the system are described in reference to block 108 of FIG. 1 of U.S. patent application Ser. No. 14/473,860, titled “MALWARE DATA ITEM ANALYSIS,” previously incorporated by reference herein. For example, academic analyses may be gathered, the suspected malware file may be executed in a sandbox environment, third-party analyses (for example FireEye and/or VirusTotal) may be run, files related to the file data item may be gathered (for example, payloads delivered in the sandbox environment), and/or the like. At block 2408 (of FIG. 24A) any information and/or data generated by the basic and/or external analysis of the file data item are associated with the file data item. Such analysis data may be referred to herein as “analysis information” and/or “analysis information items.” Examples of associating analysis data with the file data item, and examples of analysis information items, are described in reference to blocks 110 and 112 of FIG. 1 of U.S. patent application Ser. No. 14/473,860, titled “MALWARE DATA ITEM ANALYSIS,” previously incorporated by reference herein. For example, analysis information items associated with the file data item may include calculated hashes, file properties (for example, file name, file size, and/or the like), academic analysis information, file execution information (for example, effects of executing the file data item in a sandbox environment such as file system and registry changes, payloads delivered, and/or the like), third-party analysis information, and/or the like.

Turning again to FIG. 24A, at block 2410 a human analyst's evaluation of the file data item is received by the system. The analyst may provide the evaluation via, for example, a user interface of the data analysis system as described in reference to blocks 114, 120, and/or 122 of FIG. 1 of U.S. patent application Ser. No. 14/473,860, titled "MALWARE DATA ITEM ANALYSIS," previously incorporated by reference herein. For example, the analyst may review the analysis data associated with the file data item (for example, the external analysis including files produced by the file data item, actions taken by the file data item, URLs contacted by the file data item, and/or the like), determine that the file data item is likely to be malware, and may then mark the file data item as, for example, "malicious." In other embodiments the analyst may mark the file data item as likely malware in other ways suitable for the system to be notified that the evaluation resulted in a finding that the file data item is likely malware.

At block 2412, when the system determines that the file data item is not marked as "malicious" (and/or otherwise not likely malware), the method proceeds to block 2414 where it ends with respect to that submitted file data item. However, when the system determined that the file data item is marked as "malicious" (and/or otherwise likely malware), the method proceeds to block 2416 where the file data item is designated as a seed. Accordingly, the data analysis system may designate and use the file data items marked as "malicious" as seeds.

In some embodiments, one or more aspects of the basic and/or external analyses may or may not be initiated, and/or may be initiated in a different order. For example, the analysis by third-parties may not be initiated. In various embodiments, one or more aspects of the basic and/or external analyses may be performed by the data analysis system and/or other computer systems.

ii. Externally Identified Seeds

As also mentioned above, according to an embodiment seeds may be generated by the data analysis system in the context of malware threat analysis based on external threat lists. Such an embodiment is described below in reference to FIG. 25A.

Turning now to the embodiment of FIG. 25A, at block 2502 one or more external threat lists are scanned by the system. External threat lists may include one or more lists or feeds of information related to malware threats. Such threat lists may be available for free or for payment (for example, by subscription) from third-party sources. For example, such threat lists may be available from vendors of anti-malware and other security software. Threat lists may comprise, for example, blacklists provided by security services vendors. Threat lists may comprise frequently, or continuously, updated feeds of information on malware threats. Information provided by such feeds may include, for example, names of malware files, effects of or actions taken by malware files, hashes or other identifiers of malware files, information transmitted or received by malware files, domains or URLs which are contacted by or which contact malware files, and/or the like. Such threat lists are further typically accessible, hosted, and/or stored by the third parties and may be accessible via a network external to the internal network of the organization, for example, accessible via the Internet. Such threat lists are compiled and frequently updated by the third-party sources as new malware is identified for any available source.

At block 2504, the system identifies, from the scanned external threat lists, one or more external domains and/or URLs. For example, URLs (or domains extracted from URLs) on one or more threat lists that are indicated as related to malware threats (or other similar indicator) may be identi-

fied for use as seeds. For example, a domain, such as "examplemalwaredomain.com," may be identified as being contacted by a particular item of malware on the threat list. Accordingly, such a domain may be understood to be related to undesirable malware, or malicious, activity. Any such domains and/or URLs are identified and extracted from the external threat lists.

At block 2506, the system designates any identified domains and/or URLs as seeds. Accordingly, the data analysis system may designate and use the domains and/or URLs as seeds. The designated domains and/or URLs may be referred to herein as "external domain data items."

b. Malware Threat Detection: Cluster Generation

FIGS. 24B and 25B are flowcharts showing example methods of cluster generation performed by the data analysis system as applied to malware threat detection, according to embodiments of the present disclosure. The cluster generation methods of FIGS. 24B and 25B may generally be understood to correspond to block 920 (generate cluster(s)) of the generalized process of FIG. 10A. FIG. 24B relates to an embodiment of the data analysis system in which clusters are generated based on file data items designated as seeds (as described above in reference to FIG. 24A), while FIG. 25B relates to an embodiment of the data analysis system in which clusters are generated based on domains and/or URLs designated as seeds (as described above in reference to FIG. 25A).

i. Cluster Generation Based on File Data Items

As mentioned above, according to an embodiment clusters may be generated by the data analysis system in the context of malware threat analysis based on designated file data items. Such an embodiment is described below in reference to FIG. 24B.

Turning to the embodiment of FIG. 24B, at optional block 2442, one or more data stores may be scanned by the system to find any file data items marked as seeds (as described above). Block 2444 of the flowchart of FIG. 24B indicates that the following blocks (2446, 2448, 2450, and 2452) may be performed for each of the file data items marked as seeds (as described above in reference to FIG. 24A).

At block 2446, the system extracts from the file data item any network indicators. Network indicators may include, for example, domains, URLs, and/or any other indicator of network communications. For example, a particular file data item may include an indication of a connection to the domain "http://maliciousdomainexample.com" as identified by the file data item analysis described above (for example, the basic analysis of block 2404 and/or the external analysis of block 2406 described in reference to FIG. 24A above). Accordingly, the system extracts such a domain, the domain being a network indicator, from the file data item.

At block 2448, the system scans communications traffic between the internal network and the external network to identify any hosts (and/or other devices) of the internal network accessing one of the extracted network indicators. Further, the system scans for any communications to or from the extracted network indicators. Such scanning may be accomplished, for example, via a network routers, proxy devices, and/or switches as described above. Such scanning may be ongoing, meaning that network communications traffic is continuously monitored and connections are identified. Alternatively, network communication traffic may be periodically or intermittently scanned. In another embodiment, network communications traffic may be continuously (and/or periodically and/or intermittently) scanned or monitored and logged in a data store. In this embodiment, the system may continuously (and/or periodically and/or intermittently) scan the

61

stored logs of the network communications traffic to identify any hosts of the internal network accessing one of the extracted network indicators.

Hosts of the internal network may include, for example, computerized devices or systems (for example, servers, laptop computers, desktop computers, and/or the like). Hosts may be identified according to, for example, an associated computing device IDs, associated IP addresses, and/or the like). Hosts of the internal network accessing extracted network indicators are identified as potentially being infected by malware associated with the respective network indicators. Examples of hosts and monitoring of network communications are described in U.S. patent application Ser. No. 14/139,603, previously incorporated by reference herein for all purposes.

At block **2450**, data items related to the file data item marked as a seed, including the network indicators, identified hosts, and users, are clustered by the system. Clustering of data items may be accomplished as generally described above (for example, by searching various data sources for related data items and clustering those data items) as part of a clustering strategy. In general, the seed file data item is added to a cluster (for example, a data cluster is instantiated with an initial file data item being the seed or the seed is initially designated as the data cluster), and as related data items are found (for example, the identified hosts, network indicators, and/or related file data items, among others) they are added to the cluster. As described above, clustered data items may be related by, for example, sharing the same or similar properties, characteristics, and/or metadata, and are determined based on rules of one or more clustering strategies, such as a strategy that is particular to internal threat detection.

In the context of malware threat detection, data items that may be clustered with a file data item seed may be referred to as network-related data items, and may be related to the network communications described above which may be scanned and/or logged. Examples of such network-related data items include, but are not limited to: users (for example, persons having accounts on particular computer systems or devices), organizational or hierarchical positions associated with users, external Internet Protocol addresses, external domains, external IP addresses associated with external domains, external computerized devices or systems, internal Internet Protocol addresses, internal computerized devices or systems (also referred to herein as hosts), internal IP addresses that connect to external domains, internal computer systems that connect to external domains, other data items (for example, data items drawn from threat lists and/or public and/or private whitelists or blacklists, such as data items representing malware, known bad domains, known good domains, known bad IP addresses, and/or the like), host-based events (such as, for example, virus scan alerts and/or logged events, intrusion prevention system alerts and/or logged events, and the like), and/or the like.

As described above, the system may iteratively generate clusters such that previously clustered data items (and/or properties, characteristics, and/or metadata related to those clustered data items) may be used as criteria for further clustering of related data items. For example, the cluster generation method may optionally repeat multiple times until, for example, the clustering strategy is completed and/or no additional related data items are found by the system.

In general, each data cluster generated by the system in the context of malware file data item threat detection corresponds to a particular malware file (and/or group of related malware files) as represented by the file data item. For example, when a new file data item is submitted to the system and marked as

62

a seed, if that file data item has previously been submitted as a seed, that file data item may be merged into the cluster of the previously submitted file data item. Such a determination that the two file data items are the same may occur during seed generation (as described above and in reference to U.S. patent application Ser. No. 14/473,860, titled "MALWARE DATA ITEM ANALYSIS") or, alternatively, it may occur during cluster generation (as described above). In some instances a cluster may include multiple different file data items that are related by, for example, a particular domain that both contact, and/or one or more other characteristics.

As indicated in FIG. **24B**, clusters may be generated periodically as network communications traffic is continuously logged and/or scanned. For example, a cluster may be initially generated based on a seed file data item including a network indicator "http://maliciousdomainexample.com," and various communications and related hosts and users may be clustered. The cluster may then be periodically regenerated and/or updated as new communications traffic from, for example, one or more hosts may include further connections to "http://maliciousdomainexample.com." In this example, as new connections are made by hosts to network indicators associated with the cluster, those hosts and related data items are added to the cluster via, for example, cluster updating and/or regeneration. Such cluster regeneration is described above in reference to blocks **1006** and **1008** of FIG. **10A**. As described above, clusters may be regenerated on a schedule, as needed, and/or on demand.

As described above in reference to blocks **1002** and **1004** of FIG. **10A**, the system additionally analyzes generated clusters, generates alerts, and provides alerts to the analyst via a user interface. An example user interface is described below in reference to FIG. **24C**.

As indicated at block **2452** of FIG. **24B**, and as described above, as clusters are generated and regenerated, they may then be reanalyzed and alerts may be updated, linked, and/or merged.

Further examples of cluster generation and iterative clustering are described in U.S. patent application Ser. No. 14/139,603, previously incorporated by reference herein for all purposes.

ii. Cluster Generation Based on Domains/URLs

As also mentioned above, according to an embodiment clusters may be generated by the data analysis system in the context of malware threat analysis based on domains and/or URLs designated as seeds. Such an embodiment is described below in reference to FIG. **25B**.

Referring to FIG. **25B**, block **2544** of the flowchart indicates that the following blocks (**2548**, **2550**, and **2552**) may be performed for each of the external domain data items designated as seeds (as described above in reference to FIG. **25A**).

At block **2548**, the system scans communications traffic between the internal network and the external network to identify any hosts (and/or other devices) of the internal network accessing one of the external domain items. Further, the system scans for any communications to or from the external domain data items. This is generally accomplished as described above in reference to block **2448** of FIG. **24B**.

At block **2550**, data items related to the external domain data item, including the identified hosts and users, are clustered by the system. This is generally accomplished as described above in reference to block **2450** of FIG. **24B**. As described above, in the context of malware threat detection, data items that may be clustered may be referred to as network-related data items. In general, each data cluster generated by the system in the context of malware domain threat detection corresponds to a particular malware (and/or likely

malicious) domain or URL (and/or group of domains or URLs) as represented by the domain data item.

As described above, clusters may be regenerated and/or merged as new domain data items are received and/or related as the same as (or, alternatively, similar to) previously received domain data items. Further, clusters may be generated periodically as network communications traffic is continuously logged and/or scanned.

As described above in reference to blocks 1002 and 1004 of FIG. 10A, the system additionally analyzes generated clusters, generates alerts, and provides alerts to the analyst via a user interface. An example user interface is described below in reference to FIG. 24C.

As indicated at block 2552 of FIG. 25B, and as described above, as clusters are generated and regenerated, they may then be reanalyzed and alerts may be updated, linked, and/or merged.

In an embodiment, data items identified that relate to a cluster (for example, additional communications and/or host communications with a particular domain data item) are only added to the cluster if they occur within a particular time span as other clustered data items, or within a particular window of time as other clustered data items. For example, in the example of a same-day limitation, a cluster may include hosts that communicated with a particular malicious domain on July 17. When the system identifies new communications between another host and the particular malicious domain on July 17, the other host may be added to the cluster. However, when the system identifies additional communications between yet another host and the particular malicious domain on July 18, the yet other host may not be added to the cluster. Rather, a new cluster may be generated including the particular domain data item and the yet other host. Accordingly, in this embodiment, clusters are not merged when events in those clusters occur during different time periods (for example, on different days).

c. Malware Threat Detection: Alert Generation/User Interface

As mentioned above, in various embodiments the system generates alerts corresponding to each generated cluster, and provides an interface including those alerts, as described above in reference to block 1002 and 1004 of FIG. 10A. Generation of alerts, according to various embodiments, includes analysis of the generated cluster based on analysis rules/criteria specific to the type of cluster (in this example, malware threat detection), generation of conclusions, and generation of alert score(s), as described in reference to FIG. 10B above.

Examples of conclusion templates that may be applied by the system in the context of malware threat detection are shown in FIG. 10C described above. For example, "Internal Threat Intel" lists examples of conclusion templates that may be used with clusters based on file data items, as described above, while "External Threat Intel" lists examples of conclusion templates that may be used with clusters based on external domain data items, as also described above.

FIG. 24C illustrates an example data cluster analysis user interface 2472 of the data analysis system as applied to malware threat detection, according to an embodiment of the present disclosure. The analysis user interface 2472 is similar to the analysis user interface 1202 described above in reference to FIG. 12. Accordingly, the description above in reference to FIG. 12 may be understood to apply to FIG. 24C, where appropriate.

The analysis user interface 2472 includes header information 2476 including a title indicating that this alert relates to "Malicious Traffic by Malware.exe" (where "malware.exe"

may be associated with a file data item identified as a seed, for example), and was generated based on a cluster generated by the "Internal Threats" clustering strategy (for example, the seed generation and clustering strategies described above in reference to FIGS. 24A-24B). An upper portion of the analysis user interface 2474 is colored according to an alert score generated for the cluster. The analysis user interface further includes four conclusions 2478 (similar to the "Internal Threat Intel" template conclusions listed in FIG. 10C) indicating, for example, that "malware.exe" is related to three network indicators (in this example, URLs), that twelve hosts made connections to those URLs while 5 more hosts made connection to domains and/or IPs related to those URLs (and/or IPs and/or domains otherwise included in the cluster), that 5% of proxy traffic from those hosts was blocked and the last connection made by a host to one of those URLs (and/or domains and/or IPs) on Dec. 6, 2011, and that the proxy device categorized 2% of the network traffic from those hosts as "Malicious/Botnet." As mentioned above, the conclusions may further include, among other information, information related to the "bands" associated with users, for example, a highest band associated with any user in the cluster. (See, for example, the conclusion template associated with "External Threat Intel": "Highest band was <v>.") Further, the analysis user interface includes various detailed information regarding the cluster, including latest hosts 2480 in the cluster (which information may include, for example, timestamps, hostnames and/or computing device IDs associated with hosts, user accounts and IP addresses associated with those hosts at the time of the timestamp, and/or external URLs connected to at the time of the timestamp, just to name a few), and latest traffic in the cluster (similar to the description of FIGS. 12-13 above). Further, as described above, the analyst may select "Show Logs" to view even further detailed information from the cluster, as described above in reference to FIGS. 14-17, for example.

As mentioned above in reference to FIG. 12, in an embodiment, the analysis user interface 2472 may include user interface elements (such as buttons) to add the cluster data items to a graph and/or view the cluster data items in a graph or other visualization, as described in various related applications listed above and previously incorporated by reference herein. Further, the system may enable an analyst viewing a graph of data items to go to alerts representing clusters in which that data item is included.

Advantageously, as described above and according to various embodiments, the data analysis system may automatically generate seeds related to malware threats, cluster various data items related to those seeds, analyze those clusters, generate human-readable conclusions, generate alerts and alert scores, and display a user interface including various information related to the alert and the cluster to an analyst. The analyst may then efficiently evaluate the alert and determine whether the identified file data items or external domain data item and related communications represent a malware threat critical enough that it should be further investigated. Various items of information generated by the system and provided in the alert user interface assist the analyst in performing the evaluation including, for example, the alert score, the conclusions, and the easily accessible and filterable data related to the cluster.

XII. Example Application of the Data Analysis System to Network Intrusion Detection

FIGS. 26A-26C illustrate various methods and a user interface of the data analysis system as applied to network intru-

65

sion detection, according to embodiments of the present disclosure. An Intrusion Detection System (IDS) is a device or application (for example, a software application) that monitors traffic on a network of computing devices to identify potentially malicious communications or activities. For example, an IDS may be configured to detect communications initiated or received by malware (as described above), identify malicious attempts to access a network without authorization, identify malicious attempts to access a network using stolen authentication information (for example, stolen passwords), and/or the like. In general, an IDS logs identified communications and/or activities, and provides notifications or reports regarding those identified communications and/or activities. Such reports may be provided to, for example, a network administrator such that the administrator may review the logs and take action to prevent malicious action, remove malware, and/or the like. Reports provided by an IDS may include, for example, contents of a suspicious communication, a source and destination of the communication (for example, a source IP address and a destination IP address), and/or any other information related to the identified communication and/or activity. IDS's may be located at a single place in a network that is being monitored (for example, at an access point between an internal network and an external network, for example, the Internet), and/or at multiple places in the network (for example, at each computing device on the network). In general, IDS's are arranged in a network such that traffic to and from all devices on the network may be monitored. Usually IDS's are employed such that an organization may monitor traffic (and detect malicious activity) within an internal network and between the internal network and one or more external networks. Examples of internal and external network communications, and various malicious activities, are described above, for example, in reference to FIGS. 24A-24C.

Protection of an internal network from malicious activity can be a challenging task. Even when an IDS identifies potentially malicious communications, and provides a report to an administrator, the administrator often must then review network logs and attempt to decipher what is going on, what computing devices and/or addresses are involved, and whether or not further action should be taken to protect the network.

The data analysis system may be used, as described below, to automatically, efficiently, and effectively, and based on an IDS report, identify and cluster various data items related to the IDS report. The clustered data items may include, for example, internal and external network addresses (for example, domains and/or IP addresses), host devices, users, and/or other data items related to the IDS report (as described below). Such a data cluster may be analyzed by the data analysis system and an alert (including various conclusions) may be generated and presented to an analyst. The analyst may then efficiently review the alert, accurately determine whether the alert indicates a threat (and assess a level of the threat), and take action to neutralize the threat (for example, by blocking communications to or from particular network addresses (for example, addresses associated with external computing devices from which an intrusion attempt originates), cleaning computing devices infected with malware, securing computing devices that have been compromised, and/or the like. Further, as the alert/cluster includes users (for example, data items representing users) associated with computing devices in the cluster, the analyst may determine particular users and user behaviors that may have contributed to the intrusion attempt (and/or other malicious activity), and mitigate future risks (through, for example, user education).

66

As with the malware threat detection applications described above, clustered data items in this application may include organizational and/or hierarchical positions or groups (also referred to herein as "bands") associated with the clustered users.

Further, as described above in reference to the malware threat detection applications, examples of the data analysis system as applied to detection of malware (including seed generation, clustering, and cluster analysis and scoring) are described in U.S. patent application Ser. No. 14/139,603, previously incorporated by reference herein. Aspects of the examples of clustering, analysis, and scoring described in U.S. patent application Ser. No. 14/139,603 may be applied to the embodiments of the data analysis system as applied to network intrusion detection described below.

As described above and below, the data analysis system may be used in a network environment in which an internal network is in communication with an external network. Various computerized devices may be included in the internal network that may be capable of capturing, analyzing, and/or logging data traffic between the internal network and the external network including, for example, IDS's, network routers, proxy devices, and/or switches.

a. Network Intrusion Detection: Seed Generation

FIG. 26A is a flowchart of an example of a seed generation method of the data analysis system as applied to network intrusion detection, according to an embodiment of the present disclosure. The seed generation method of FIG. 26A may generally be understood to correspond to block 910 (generate seed(s)) of the generalized process of FIG. 10A.

At block 2602, one or more Intrusion Detection System (IDS) reports are received by the system. In the embodiment of FIG. 26A, the IDS report(s) are received from one or more IDS devices located within the internal network, and may be provided by third parties. In another embodiment, the IDS(s) may be integrated with the data analysis system and/or may comprise one or more software applications installed on computing devices of the internal network (for example, servers, desktop computers, laptop computers, routers, proxy devices, and/or the like).

At block 2604, the received IDS reports may optionally be stored by the system in, for example, one or more data stores as described above. Then, at block 2606, the IDS reports may be scanned by the system to extract and/or parse information from the IDS reports. Alternatively, the IDS reports may be scanned by the system immediately and/or shortly after being received, and the extracted and/or parsed information may then be used by the system (as described below) and/or temporarily stored in, for example, one or more data stores. In various embodiments IDS reports may be scanned on a schedule, as needed, and/or on demand (similar to the description of blocks 1006a, 1006b, and 1006c of FIG. 10A given above). For example, in an embodiment received IDS reports are stored as they are received, and then periodically scanned. In another embodiment, IDS reports may be scanned continuously (for example, immediately or substantially immediately upon receipt) by the system.

Information extracted and/or parsed from the IDS reports may include, for example, contents of (for example, data associated with) a suspicious communication, a source and destination of the communication (for example, a source IP address and a destination IP address), and/or the like.

At block 2608, the system performs a WHOIS lookup on the extracted source and destination IP addresses (and/or, for example, source and destination domain addresses) associated with each IDS report so as to identify which address is internal to the internal network and with address is external to

67

the internal network. In general, a WHOIS lookup comprises a query of a database of information related to various IP addresses and/or domain addresses. The information provided in response to a WHOIS lookup may include, for example, a location of a computing device associated with the address, registered users associated with the address, and/or the like. In an embodiment, the WHOIS lookup may be directed to a third-party WHOIS service provider. In another embodiment, a WHOIS lookup service and/or database may be maintained by the data analysis system.

Based on the response to the WHOIS lookup, the system determines which of the addresses associated with a particular IDS report are external and which are internal. For example, the WHOIS lookup may indicate that a computing device associated with a particular IP address is located in a country (and/or address) different from any in which the organization's network is located.

In another embodiment, the system may access to a list of IP addresses associated with the organization. In this embodiment, the system may compare IP addresses associated with a particular IDS report to the list of IP addresses associated with the organization to determine which IP addresses associated with the particular IDS report are external and which are internal. Information regarding IP addresses associated with, or assigned to, and organization may be accessed by the system from, for example, the Internet Assigned Numbers Authority, one or more Regional Internet registries, and/or any other provider of IP address allocation and/or assignment information provider.

Accordingly, in any of the embodiments described above, the system may determine which IP addresses are external to the internal network. Such addresses are referred to herein as "external addresses" and, at block 2610, the system designates and uses those external addresses as seeds.

b. Network Intrusion Detection: Cluster Generation

FIG. 26B is a flowchart of an example of a clustering method performed by the data analysis system as applied to network intrusion detection, according to an embodiment of the present disclosure. The cluster generation method of FIG. 26B may generally be understood to correspond to block 920 (generate clusters(s)) of the generalized process of FIG. 10A. Block 2644 of the flowchart of FIG. 26B indicates that the following blocks (2646 and 2650) may be performed for each of the external addresses marked as seeds (as described above in reference to FIG. 26A).

At block 2646, the system clusters any data items that are related to the seed including, for example, internal addresses extracted from the IDS report (for example, internal IP addresses and/or domain addresses), any other information extracted from the IDS report, any information gathered via the WHOIS lookups, hosts associated with the network addresses, users associated with the hosts, and/or the like. Additionally, other network traffic information clustered may include information gathered from firewall devices and/or routers of the network (for example, fingerprints, signatures, and/or hashes associated with malware items and/or particular communications).

Clustering of data items may be accomplished as generally described above (for example, by searching various data sources for related data items and clustering those data items) as part of a clustering strategy. In general, the seed file data item is added to a cluster (for example, a data cluster is instantiated with an initial clustered data item being the seed or the seed is initially designated as the data cluster), and as related data items are found (for example, the various data items mentioned above and below) they are added to the cluster. As described above, clustered data items may be

68

related by, for example, sharing the same or similar properties, characteristics, and/or metadata.

In the context of network intrusion detection, data items that may be clustered may be referred to as network-related data items, and may be related to the network communications described above which may be scanned and/or logged. Examples of such network-related data items include, but are not limited to: users (for example, persons having accounts on particular computer systems or devices), organizational or hierarchical positions associated with users (referred to herein as "band" levels and described above), external Internet Protocol addresses, external domains, external IP addresses associated with external domains, external computerized devices or systems, internal Internet Protocol addresses, internal computerized devices or systems (also referred to herein as hosts), internal IP addresses that connect to external domains, internal computer systems that connect to external domains, other data items (for example, data items drawn from threat lists and/or public and/or private whitelists or blacklists, such as data items representing malware, known bad domains, known good domains, known bad IP addresses, and/or the like), host-based events (such as, for example, virus scan alerts and/or logged events, intrusion prevention system alerts and/or logged events, and the like), WHOIS information associated with network addresses (for example, geographical location information and/or registration information), information extracted from IDS's (as described above), network firewall- and/or router-gathered information (for example, fingerprints, signatures, and/or hashes associated with malware items and/or particular communications), and/or the like.

As described above, and as indicated by decision block 2648, the system may iteratively generate clusters such that previously clustered data items (and/or properties, characteristics, and/or metadata related to those clustered data items) may be used as criteria for further clustering of related data items. For example, the cluster generation method may optionally repeat multiple times until, for example, the clustering strategy is completed and/or no additional related data items are found by the system. When the clustering strategy is complete, the method may continue to block 2650, as described below.

In general, each data cluster generated by the system in the context of network intrusion detection corresponds to a particular external address if the corresponding IDS reports were received within a common time period (as described below). For example, when a new IDS report is received by the system, and a previously identified external address is re-identified in the new IDS report, that external address may be merged into the cluster of the previously identified external address (in an embodiment, subject to the time-period limitations described below). Such a determination that the two identified external addresses are the same may occur during seed generation or, alternatively, it may occur during cluster generation.

As indicated by block 2650, and as mentioned above, in an embodiment clusters of data items based on common external addresses may only be merged if the corresponding IDS reports were received within a particular time period, time span, or window of time (for example, on a same day, in a same hour, and/or the like). For example, in the example of a same-day limitation, a cluster may be based on an external address that was extracted from an IDS report received on July 17. When the system identifies that same external address in a later received IDS report, if the later received IDS report is also received on July 17, the external address (and any related data items) may be merged into the same cluster.

However, if the later received IDS report is received on July 18, a new cluster may be generated by the system and the external address (and any related data items) may be added to that new cluster. In another embodiment, no time limitation is placed on cluster merging.

As mentioned above, seeds may be identified intermittently (for example, on a schedule, as needed, and/or on demand) and/or continuously as IDS reports are received. Similarly, in various embodiments clusters and/or alerts (as described above and below) may be generated, merged, and/or regenerated intermittently and/or continuously as seeds are identified/generated.

As described above in reference to blocks **1002** and **1004** of FIG. **10A**, the system analyzes generated clusters, generates alerts, and provides alerts to the analyst via a user interface. An example user interface is described below in reference to FIG. **26C**.

In general, alerts/clusters are generated/regenerated and analyzed/reanalyzed as described above such that alerts are may be updated, linked, and/or merged.

Further examples of cluster generation and iterative clustering are described in U.S. patent application Ser. No. 14/139,603, previously incorporated by reference herein for all purposes.

c. Network Intrusion Detection: Alert Generation/User Interface

As mentioned above, in various embodiments the system generates alerts corresponding to each generated cluster, and provides an interface including those alerts, as described above in reference to block **1002** and **1004** of FIG. **10A**. Generation of alerts, according to various embodiments, includes analysis of the generated cluster based on analysis rules/criteria specific to the type of cluster (in this example, malware threat detection), generation of conclusions, and generation of alert score(s), as described in reference to FIG. **10B** above.

Examples of conclusion templates that may be applied by the system in the context of network intrusion detection are shown in FIG. **10C** described above. For example, “IDS” lists examples of conclusion templates that may be used with clusters based on external addresses, as described above.

FIG. **26C** illustrates an example data cluster analysis user interface **2672** of the data analysis system as applied to network intrusion detection, according to an embodiment of the present disclosure. The analysis user interface **2672** is similar to the analysis user interface **1202** described above in reference to FIG. **12**. Accordingly, the description above in reference to FIG. **12** may be understood to apply to FIG. **26C**, where appropriate.

The analysis user interface **2672** includes header information **2676** including a title indicating that this alert relates to “10 Outbound/Inbound IDS Reports To/From 7.23.0.1” (where “7.23.0.1” may be an external IP address identified as a seed, for example), and was generated based on a cluster generated by the “IDS” clustering strategy (for example, the seed generation and clustering strategies described above in reference to FIGS. **26A-26B**). An upper portion of the analysis user interface **2674** is colored according to an alert score generated for the cluster. For example, in one embodiment the upper portion is colored yellow to indicate a medium alert, whereas a high alert may be associated with red coloring, and a low alert may be associated with green coloring (or any other coloring, formatting, or display settings). The analysis user interface further includes four conclusions **2678** (similar to the “IDS” template conclusions listed in FIG. **10C**) indicating, for example, that the computing device associated with the external IP address is registered to Secret, Inc. based

in China; that the external IP address triggered 10 IDS reports related to 3 different host computing devices (on the internal network); that the received IDS reports in the cluster spanned 10 hours, 5 minutes, starting at 12:35; and that the most common signature associated with the traffic captures by the IDS was EXAMPLE SIGNATURE. As mentioned above, the conclusions may further include, among other information, information related to the users of the identified hosts and/or “bands” associated with the users, for example, a highest band associated with any user in the cluster. Further, the analysis user interface includes various detailed information regarding the cluster, including latest IDS Report Reference Numbers (for example, reference numbers associated with most recently received IDS reports that are included in the cluster) in the cluster. Further, as described above, the analyst may select “Show Logs” to view even further detailed information from the cluster, as described above in reference to FIGS. **14-17**, for example.

As mentioned above in reference to FIG. **12**, in an embodiment, the analysis user interface **2672** may include user interface elements (such as buttons) to add the cluster data items to a graph and/or to display the cluster data items in a graph or any other visualization, as described in various related applications listed above and previously incorporated by reference herein. Further, the system may enable an analyst viewing a graph of data items to go to alerts representing clusters in which that data item is included.

Advantageously, as described above and according to various embodiments, the data analysis system may automatically generate seeds related to potential malicious network activity (for example, intrusion attempts), cluster various data items related to those seeds, analyze those clusters, generate human-readable conclusions, generate alerts and alert scores, and display a user interface including the various information related to the alert and the cluster to an analyst. The analyst may then efficiently evaluate the alert and determine whether the identified malicious activities and related data items represent an intrusion threat (or other malicious threat) critical enough that it should be further investigated. Various items of information generated by the system and provided in the alert user interface assist the analyst in performing the evaluation including, for example, the alert score, the conclusions, and the easily accessible and filterable data related to the cluster.

XIII. Example Application of the Data Analysis System to Phishing Detection

FIGS. **27A-27C** illustrate various methods and a user interface of the data analysis system as applied to phishing threat detection, according to embodiments of the present disclosure.

The term “phishing,” as used herein, is a broad term encompassing its plain and ordinary meaning and includes, without limitation, all types of attempts to acquire sensitive information (for example, private information, usernames, passwords, credit card or other payment information, and/or the like) from a target by false representation in a communication (especially, for example, in an electronic communication such as an email). For example, a malicious party (for example, an individual, group of individuals, and/or organization) may send a phishing email to a customer of a bank falsely representing that the phishing email is from the bank. The phishing email may request that the customer click a link or reply to the email with personal details, however in doing so that customer would be disclosing sensitive information not to the bank, but to the malicious party.

71

Phishing “campaigns” may comprise coordinated and/or large-scale efforts to dupe one or more targets by sending multiple communications (for example, emails) to the one or more targets. For example, many employees and/or customers of a particular organization, such as a bank, may all be sent similar phishing emails as part of a phishing campaign.

Legitimate organizations, such as business, seeking to establish the trust of their customers and to prevent disclosure of sensitive information by their customers and/or their employees to malicious parties have a strong interest in detecting and combating phishing attempts and threats (including phishing campaigns). Previous tools for combating phishing have been limited at least in that much of the information needed to assess the seriousness and extent of a phishing effort were not accessible, or accessing such information was difficult and time consuming.

Phishing attempts may threaten an organization on at least two fronts: 1. phishing may target employees of the organization in an effort to get the employees to disclose sensitive information about the organization (for example, usernames, passwords, and/or the like) that may give a malicious party access to even further sensitive information of the organization and/or may allow the malicious party to access and manipulate or destroy the organization’s data; and 2. phishing may target customers of the organization in an effort to get the customers to disclose sensitive information about themselves (for example, usernames, passwords, credit card or other payment information, and/or the like). Described below are applications of the data analysis system to detecting phishing efforts on both of the fronts mentioned above, according to various embodiments.

In particular, the data analysis system may be used, as described below, to automatically, efficiently, and effectively detect phishing threats both internal and external to an organization, and identify and cluster various data items related to the detected phishing threats. The clustered data items may include, for example, related phishing emails received by multiple individuals, links (for example, Uniform Resource Locators (URLs) and/or the like) included in those emails, and/or other data items. Such a data cluster may be analyzed by the data analysis system and an alert (including various conclusions) may be generated and presented to an analyst. The analyst may then efficiently review the alert, accurately determine whether the alert indicates a phishing threat (and assess a level of the threat), and take action to combat the phishing threat. Further, as the alert/cluster includes individuals (for example, data items representing users and/or customers) associated with the phishing efforts (for example, that received phishing emails and/or clicked on links in phishing emails), the analyst may determine particular individuals and individual behaviors that may have contributed to the phishing threat, and mitigate risks (through, for example, user education, warning individuals about the threat, preventing access to sensitive information of affected individuals, and/or the like). As with the malware threat detection applications described above, clustered data items in this application may include organizational and/or hierarchical positions or groups (also referred to herein as “bands”) associated with the clustered individuals.

Further, as described above in reference to the malware threat detection applications, examples of the data analysis system (including seed generation, clustering, and cluster analysis and scoring) in various applications are described in U.S. patent application Ser. No. 14/139,603, previously incorporated by reference herein. Aspects of the examples of clustering, analysis, and scoring described in U.S. patent application Ser. No. 14/139,603 may be applied to the

72

embodiments of the data analysis system as applied to phishing threat detection described below.

In the context of the presently described embodiments, the data analysis system may be used in a communications network environment of an organization. Various devices and software applications may be included in the communications network that may be capable of capturing and/or logging data traffic between various devices of the communications network and/or other computerized devices outside of the network including, for example, network routers, proxy devices, and/or switches.

a. Internal Phishing Detection: Seed Generation & Cluster Generation

FIG. 27A is a flowchart of an example of a seed generation and clustering method performed by the data analysis system as applied to phishing threat detection, according to an embodiment of the present disclosure. The embodiment of FIG. 27A applies to detection of phishing threats to employees of an organization, referred to herein as “Internal Phishing Detection.” The seed generation method of block 2702 (described below) may generally be understood to correspond to block 910 (generate seed(s)) of the generalized process of FIG. 10A, while the cluster generation method of blocks 2704, 2706, 2708, 2710, 2712, and 2714 may generally be understood to correspond to block 920 (generate clusters(s)) of the generalized process of FIG. 10A.

At block 2702, the system receives copies of one or more potentially malicious (for example, phishing-related) emails from employees of the organization. The copies of the emails may be sent to a particular email address, or placed in a particular electronic mailbox, of the organization designated for receiving potentially malicious emails targeting the employees of the organization. For example, employees may have a user interface control (e.g., a “report spam/phishing” button) as part of their email system that they can select in order to report a particular email as a potential phishing threat. In another embodiment, the system may automatically scan electronic mailboxes of employees to identify potential phishing emails. The system may then designate and use each of the potentially malicious emails (whether received from employees, scanned for, or otherwise obtained by the system) as seeds. Emails in the context of the present disclosure may be referred to as email data items.

Block 2704 indicates that the following blocks (2706, 2708, 2710, 2712, 2714, and 2716) may be performed for each of the seeds (for example, email data items) designated as described above in reference to block 2702.

At block 2706, the system identifies and extracts the subject and sender information from the email data item designated as the seed. The subject information of the phishing email may be, for example, “Click this link!” The sender information of the phishing email may include, for example, a source email address such as “evil@examplephishingdomain.com,” and may also include a name associated with the source email address, such as “John Doe.”

At block 2706, the system clusters any other email data items that are related to the seed based on the extracted subject information or sender information. For example, the system may cluster other email data items having a same (or, alternatively, substantially the same or similar) subject information as the seed email data item. Further, the system may cluster other email data items having a same (or, alternatively, substantially the same or similar) sender information as the seed email data item.

Clustering of data items may be accomplished as generally described above (for example, by searching various data

sources for related data items and clustering those data items) as part of a clustering strategy. In general, the seed email data item is added to a cluster (for example, a data cluster is instantiated with an initial clustered data item being the seed or the seed is initially designated as the data cluster), and as related data items are found (for example, the email data items mentioned above and below) they are added to the cluster. As described above, clustered data items may be related by, for example, sharing the same or similar properties, characteristics, and/or metadata (for example, subject and/or sender information). Email data items to be clustered may be stored in, for example, one or more data stores of the system or another computing system. In an embodiment, the system may access electronic mailboxes of the employees of the organization to obtain email data items to be clustered by the system. In an embodiment the email data items may be obtained by the system on an ongoing or continuous basis and the emails are transmitted through the communications network.

In an embodiment, the system may require (e.g., based on the Internal Phishing clustering strategy) that either the subject information or the sender information of an email data item be the same or substantially the same as the seed email data item for the email data item to be clustered. In another embodiment, the system may require that both the subject information and the sender information of an email data item be the same or substantially the same as the seed email data item for the email data item to be clustered. In yet another embodiment, the system may operate on a spectrum that determines a closeness of a match between either the subject information or the sender information of an email data item and the same as the seed email data item to determine whether the email data item is to be clustered. For example, in one instance the subject information may be completely different, while the sender information may be very nearly the same. In this example, the closeness of the match of both the sender and subject would be weighed and the system may determine that the email data item should be clustered. In another example, the subject information may be very nearly the same, while the sender information may be completely different. In this example, the closeness of the match of both the sender and subject would be weighed and the system may determine that the email data item should be clustered. In yet another example, the subject information may be mostly different, while the sender information may also be mostly different. In this example, the closeness of the match of both the sender and subject would be weighed and the system may determine that the email data item should not be clustered.

In various embodiments portions of the sender information and/or portions of the subject information may be used to determine clustering, and various items may be weighted differently in the determination. For example, an email data item has a sender email address of "john@examplephishingdomain.com," while the seed data item has a sender email address of "sam@examplephishingdomain.com." In this example, while the sender email addresses are different, the system may determine that because the domains of the two email addresses are the same, the email data item should be clustered with the seed email data item.

In various embodiments, two email data items having subject or sender information that is "substantially the same" or "similar" to one another may be understood to mean that the information need not be identical, but at least one or more portions of the information should be identical in some respect (for example, characters, structure, formatting, and/or the like). For example, two subjects may be "substantially the

same" when both include at least a portion of characters that are the same, such as "Click this link!" and "Click this link now!" In another example, subjects may be "substantially the same" when both are structured similarly, such as "Click this button!" and "Select this link!" The system may employ one or more algorithms, such as machine learning algorithms to determine whether the email data items match as described above.

In general, each data cluster generated by the system in the context of phishing detection may be considered to correspond to a particular phishing campaign. This is because, for example, all clustered email data items may be similar to one another (for example, similar subject and/or sender), and may be considered to be part of a coordinated and/or large-scale phishing effort. Accordingly, when, for example, a new email data item is received by the system that includes sender or subject information similar to a clustered email data item, that new email data item may be clustered/merged into the existing cluster. Such a determination that the two email data items are to be clustered may occur during seed generation or, alternatively, it may occur during cluster generation.

Phishing campaigns may be time dependent, for example, many emails associated with a campaign may be sent out at once (or within a short period of time). Accordingly, as indicated in block 2708, in an embodiment clustered email data items may be limited to a particular time period in addition to having similar subjects or senders. For example, when a new email data item is received by the system that includes sender or subject information similar to a clustered email data item, that new email data item may, in an embodiment, only be clustered/merged into the existing cluster when, for example, the two email data items were received either by the system or by the original recipient, within a particular time period. In various embodiments, the email data items may or may not be clustered depending on being received, either by the system or by the original recipient, within a particular time period, time span, or window of time (for example, on a same day, in a same hour, and/or the like). For example, in the example of a same-day limitation, a cluster may be based on seed email data item that was received by its recipient on July 17. When the system later identifies a similar email data item, if the later email data item was also received by its recipient on July 17, the email data item (and any related data items) may be merged into the same cluster as the seed. However, if the later email data item was received by its recipient on July 18, a new cluster may be generated by the system and the email data item (and any related data items) may be added to that new cluster. In another embodiment, no time limitation is placed on cluster merging.

At block 2710, any subject or sender information associated with any clustered email data items may be identified and added to the matching criteria for adding email data items to the cluster, and, as indicated by decision block 2716, additional email data items may then be clustered based on those new matching criteria. Accordingly, in the embodiment shown in FIG. 27A, the system may iteratively generate clusters such that previously clustered data items (and/or properties, characteristics, and/or metadata related to those clustered data items) may be used as criteria for further clustering of related data items. For example, the cluster generation method shown in block 2708 and 2710 may optionally repeat multiple times until, for example, the clustering strategy is completed and/or no additional related data items are found by the system (as indicated by decision block 2716).

In another embodiment, the method of FIG. 27A may not include block 2710 such that no additional clustering criteria

75

are added. In another embodiment, the system may only cluster email data items received from employees.

Continuing with the method illustrated in FIG. 27A, at block 2712 the system parses and/or extracts any links, for example, URLs, from the clustered email data items. These URLs may then optionally be added to the cluster as additional data items. Such URLs are typically provided by the malicious party in hopes that the target will click the URL and visit, usually, a webpage on the Internet fraudulently prepared to look like a legitimate webpage where the target should disclose sensitive information. In another example, such URLs may link to a malware file that may be downloaded to the target's computing device when selected. In either case, the system extracts such URLs, and monitors communications network traffic (as described above) so as to determine whether any phishing email recipients have clicked any of the URLs. Parsing and clustering of URLs may continue iteratively according to the clustering strategy, as described above and as indicated by decision block 2716.

At block 2714, any such employees that clicked on any of the parsed URLs are identified by the system. Such employees or users may be referred to herein to as "clickers." Clickers are identified by the system by comparing the parsed out URLs to communications network traffic. Traffic to any of the URLs may be identified by the system, a source of the traffic (for example, a network address and computing device associated with that address) may be identified by the system, users associated with that computing device (for example, an employee that was logged on to the computing device when the traffic was detected) are identified by the system, and/or organizational information associated with the users (for example, one or more "bands," as described above) may be determined by the system. This various information may be referred to, in the context of the phishing detection, as phishing-related information, and data items corresponding to this information identified and determined by the system may be clustered by the system.

In the context of phishing detection, data items that may be clustered may be referred to as phishing-related data items, and may be related to the network communications described above which may be scanned and/or logged. Examples of such phishing-related data items include, but are not limited to: users (for example, persons having accounts on particular computer systems or devices), organizational or hierarchical positions associated with users (referred to herein as "band" levels and described above), external Internet Protocol addresses, external domains, external IP addresses associated with external domains, external computerized devices or systems, internal Internet Protocol addresses, internal computerized devices or systems (also referred to herein as hosts), internal IP addresses that connect to external domains, internal computer systems that connect to external domains, other data items (for example, data items drawn from threat lists and/or public and/or private whitelists or blacklists, such as data items representing malware, known bad domains, known good domains, known bad IP addresses, and/or the like), host-based events (such as, for example, virus scan alerts and/or logged events, intrusion prevention system alerts and/or logged events, and the like), WHOIS information associated with network addresses (for example, geographical location information and/or registration information), network firewall- and/or router-gathered information (for example, fingerprints, signatures, and/or hashes associated with malware items and/or particular communications), and/or the like.

As indicated by decision block 2716, clustering may continue iteratively as new data items are added to the data

76

cluster, as described above and throughout the present disclosure. Further examples of cluster generation and iterative clustering are described in U.S. patent application Ser. No. 14/139,603, previously incorporated by reference herein for all purposes.

In various embodiments, email data item seeds may be identified intermittently (for example, on a schedule, as needed, and/or on demand) and/or continuously as copies of potentially malicious phishing emails are received by the system. Similarly, in various embodiments clusters and/or alerts (as described above and below) may be generated, merged, and/or regenerated intermittently and/or continuously as seeds are identified/generated.

As described above in reference to blocks 1002 and 1004 of FIG. 10A, the system analyzes generated clusters, generates alerts, and provides alerts to the analyst via a user interface. An example user interface is described below in reference to FIG. 27C.

In general, alerts/clusters are generated/regenerated and analyzed/reanalyzed as described above such that alerts are may be updated, linked, and/or merged.

b. External Phishing Detection: Seed Generation & Cluster Generation

FIG. 27B is a flowchart of an example of another seed generation and clustering method performed by the data analysis system as applied to phishing threat detection, according to an embodiment of the present disclosure. The embodiment of FIG. 27B applies to detection of phishing threats to customers of an organization from outside of the organization, referred to herein as "External Phishing Detection." The seed generation method of block 2742 (described below) may generally be understood to correspond to block 910 (generate seed(s)) of the generalized process of FIG. 10A, while the cluster generation method of blocks 2744, 2746, 2748, 2750, 2752, 2756, 2760, and 2762 may generally be understood to correspond to block 920 (generate clusters(s)) of the generalized process of FIG. 10A.

At block 2742, the system receives copies of one or more potentially malicious (for example, phishing-related) emails from customers of the organization. The copies of the emails may be sent to a particular email address, or placed in a particular electronic mailbox, of the organization designated for receiving potentially malicious emails targeting the customers of the organization. The system may then designate and use each of the potentially malicious emails as seeds. Emails in the context of the present disclosure may be referred to as email data items.

Block 2744 indicates that the following blocks (2744, 2746, 2748, 2750, 2752, 2756, 2760, and 2762) may be performed for each of the seeds (for example, email data items) designated as described above in reference to block 2742.

Blocks 2744, 2746, 2748, 2750, 2752, and 2756 proceed generally as described above in reference to blocks 2704, 2706, 2708, 2710, 2712, and 2716 of FIG. 27A. For example, for each of the seed email data items (block 2744), subject and sender information is identified (block 2746), email data items are clustered (blocks 2748 and 2750), URLs are parsed from the email data items and clustered (block 2752), and the process continues iteratively (decision block 2756) as described above. In contrast with Internal Phishing Detection described above, generally the system does not have access to customers' electronic mailboxes for obtaining and clustering additional email data items. In an embodiment, the system may cluster email data items accessible from third-party services, for example, services that provide spam and phishing email protection. In an embodiment, extracted URLs may be compared to identified malicious URLs provided by such

77

third-party security services (including various blacklists, for example), as described above in reference to the various applications of the data analysis system. Additional information related to the cluster may be obtained from such services and added to the data cluster.

As described above, each of the clusters generated by the system in the context of phishing detection may generally relate to a phishing campaign. Accordingly, in an embodiment, clustering/merging of email data items may be limited according to time, as described above in reference to FIG. 27A.

In various embodiments, any other phishing-related data items, as described above, may be clustered.

At block 2760, the system designates any clusters having more than a threshold number of email data items as a campaign (for example, a phishing campaign). For example, any clusters having greater than two, three, four, or five (or, alternatively, any other number of) email data items may be designated as a campaign. Clusters designated as campaigns may be referred to herein as “campaign clusters.”

At block 2762, further action may be taken by the system to interact with third-party services regarding the campaign clusters and/or information associated with the campaign clusters, and/or to automatically generate and/or analyze those clusters designated as campaigns. In an embodiment, other phishing-related data items are only clustered with clusters designated as campaigns. In another embodiment, the system only analyzes and generates alerts for clusters designated as campaigns. In yet another embodiment, whether or not a cluster is designated as a campaign is a factor weighed in determining the alert score. In another embodiment, a number of email data items in a cluster may be a factor in determining an alert score.

As mentioned, further action taken by the system may include interaction with one or more third-party services. For example, the system may determine whether or not any URLs/domains identified by the system were previously identified on one or more blacklists provided by third-party services. The system may accomplish this by comparing the identified/parsed URLs and/or domains to any URLs/domains provided on the one or more blacklists. In an embodiment the system notifies the third-party services of any URLs/domains not previously identified. In an embodiment, the third-party services are only notified regarding URLs/domains identified in relation to any campaign clusters.

In various embodiments, email data item seeds may be identified intermittently (for example, on a schedule, as needed, and/or on demand) and/or continuously as copies of potentially malicious phishing emails are received by the system. Similarly, in various embodiments clusters and/or alerts (as described above and below) may be generated, merged, and/or regenerated intermittently and/or continuously as seeds are identified/generated.

As described above in reference to blocks 1002 and 1004 of FIG. 10A, the system analyzes generated clusters, generates alerts, and provides alerts to the analyst via a user interface. An example user interface is described below in reference to FIG. 27C.

In general, alerts/clusters are generated/regenerated and analyzed/reanalyzed as described above such that alerts are may be updated, linked, and/or merged.

c. Phishing Detection: Alert Generation/User Interface

As mentioned above, in various embodiments the system generates alerts corresponding to each generated cluster, and provides an interface including those alerts, as described above in reference to block 1002 and 1004 of FIG. 10A. Generation of alerts, according to various embodiments,

78

includes analysis of the generated cluster based on analysis rules/criteria specific to the type of cluster (in this example, malware threat detection), generation of conclusions, and generation of alert score(s), as described in reference to FIG. 10B above.

Examples of conclusion templates that may be applied by the system in the context of phishing detection are shown in FIG. 10C described above. For example, the “Internal Phishing” and “External Phishing” lists show examples of conclusion templates that may be used with clusters based on email data items, as described above.

FIG. 27C illustrates an example data cluster analysis user interface of the data analysis system as applied to phishing threat detection, according to an embodiment of the present disclosure. The analysis user interface 2772 is similar to the analysis user interface 1202 described above in reference to FIG. 12. Accordingly, the description above in reference to FIG. 12 may be understood to apply to FIG. 27C, where appropriate.

The analysis user interface 2772 includes header information 2776 including a title indicating that this alert relates to “Phishing Campaign 12”, and was generated based on a cluster generated by the “Internal Phishing” clustering strategy (for example, the seed generation and clustering strategies described above in reference to FIG. 27A). An upper portion of the analysis user interface 2774 is colored according to an alert score generated for the cluster. The analysis user interface further includes four conclusions 2778 (similar to the “Internal Phishing” template conclusions listed in FIG. 10C) indicating, for example, that 9 senders (for example, malicious parties sending phishing emails) sent emails that were reported to the system (for example, to an email address such as “Abuse@bank.com”) having subjects similar to “OPEN THIS LINK!”; that the senders of those reported emails were the senders of 55 emails to the organization between June 12 and June 13; that 45 employees received those emails and a highest band of those employees was Senior Management; and that 2 domains (or, alternatively, URLs) were extracted from those emails and 12 likely clickers were identified. Further, the analysis user interface includes various detailed information regarding the cluster, including latest emails received in the cluster and a list of the latest clicks on the URLs (including who the clickers are). Further, as described above, the analyst may select “Show Logs” to view even further detailed information from the cluster, as described above in reference to FIGS. 14-17, for example.

As mentioned above in reference to FIG. 12, in an embodiment, the analysis user interface 2772 may include user interface elements (such as buttons) to add the cluster data items of a graph, as described in various related applications listed above and previously incorporated by reference herein. Further, the system may enable an analyst viewing a graph of data items to go to alerts representing clusters in which that data item is included.

Advantageously, as described above and according to various embodiments, the data analysis system may automatically generate seeds related to phishing activities (for example, malicious emails and/or phishing campaigns), cluster various data items related to those seeds, analyze those clusters, generate human-readable conclusions, generate alerts and alert scores, and display a user interface including the various information related to the alert and the cluster to an analyst. The analyst may then efficiently evaluate the alert and determine whether the identified malicious activities and related data items represent a phishing threat (or other malicious threat) critical enough that it should be further investigated. For example, the analysis user interface of FIG. 27C

indicates a fairly involved phishing campaign including a significant number of employee clickers and recipients all the way up to the Senior Management level. Accordingly, an analyst is likely to quickly recognize that this phishing campaign poses a serious threat to the organization and steps should be taken to mitigate as immediately as possible. Further, the analyst may have, in the example of FIG. 27C, all the information immediately available to determine the proper mitigation steps including, for example, domains and URLs to block, the types of incoming emails to block, and employees to educate regarding phishing threats. Various items of information generated by the system and provided in the alert user interface assist the analyst in performing the evaluation including, for example, the alert score, the conclusions, and the easily accessible and filterable data related to the cluster.

ADDITIONAL EMBODIMENTS

Embodiments of the present disclosure have been described that relate to automatic generation of memory-efficient clustered data structures and, more specifically, to automatic selection of an initial data entity of interest, adding of the initial data entity to the memory-efficient clustered data structure, determining and adding one or more related data items to the cluster, analyzing the cluster based on one or more rules or criteria to generate human-readable conclusions, and providing an interactive user interface to an analyst. As described above, in various embodiments, a generated cluster may include far fewer data items as compared to a huge collection of data items that may or may not be related to one another. This may be because, for example, data items included in a cluster may only include those data items that are related to one another and which may be relevant to a particular investigation. Accordingly, in various embodiments, processing of generated clusters may be highly efficient because, for example, a given fraud investigation by an analyst may only require storage in memory of a single cluster data structure. Further, a number of data items in a cluster may be several orders of magnitude smaller than in the huge collection of data items that may or may not be related to one another because only data items related to each other are included in the cluster.

Additionally, the automated analysis and scoring of clusters (as mentioned above) may enable highly efficient evaluation of the various data clusters by a human analyst. For example, the interactive user interface is generated so as to enable an analyst to quickly view critical data clusters (as determined by the automated scoring), and then in response to analyst inputs, view and interact with the generated information (including, for example, the human-readable conclusions) associated with the clusters. In response to user inputs the user interface may be updated to display raw data associated with each of the generated clusters if the analyst desires to dive deeper into data associated with a given cluster.

While the foregoing is directed to various embodiments, other and further embodiments may be devised without departing from the basic scope thereof. For example, aspects of the present disclosure may be implemented in hardware or software or in a combination of hardware and software. An embodiment of the disclosure may be implemented as a program product for use with a computer system. The program(s) of the program product define functions of the embodiments (including the methods described herein) and may be contained on a variety of computer-readable storage media. Illustrative computer-readable storage media include, but are not limited to: (i) non-writable storage media (e.g., read-only memory devices within a computer such as CD-ROM disks

readable by a CD-ROM drive, flash memory, ROM chips or any type of solid-state non-volatile semiconductor memory) on which information is permanently stored; and (ii) writable storage media (e.g., hard-disk drive or any type of solid-state random-access semiconductor memory) on which alterable information is stored. Each of the processes, methods, and algorithms described in the preceding sections may be embodied in, and fully or partially automated by, code modules executed by one or more computer systems or computer processors comprising computer hardware. The processes and algorithms may alternatively be implemented partially or wholly in application-specific circuitry.

The various features and processes described above may be used independently of one another, or may be combined in various ways. All possible combinations and subcombinations are intended to fall within the scope of this disclosure. In addition, certain method or process blocks may be omitted in some implementations. The methods and processes described herein are also not limited to any particular sequence, and the blocks or states relating thereto can be performed in other sequences that are appropriate. For example, described blocks or states may be performed in an order other than that specifically disclosed, or multiple blocks or states may be combined in a single block or state. The example blocks or states may be performed in serial, in parallel, or in some other manner. Blocks or states may be added to or removed from the disclosed example embodiments. The example systems and components described herein may be configured differently than described. For example, elements may be added to, removed from, or rearranged compared to the disclosed example embodiments.

Conditional language, such as, among others, “can,” “could,” “might,” or “may,” unless specifically stated otherwise, or otherwise understood within the context as used, is generally intended to convey that certain embodiments include, while other embodiments do not include, certain features, elements and/or steps. Thus, such conditional language is not generally intended to imply that features, elements and/or steps are in any way required for one or more embodiments or that one or more embodiments necessarily include logic for deciding, with or without user input or prompting, whether these features, elements and/or steps are included or are to be performed in any particular embodiment.

The term “comprising” as used herein should be given an inclusive rather than exclusive interpretation. For example, a general purpose computer comprising one or more processors should not be interpreted as excluding other computer components, and may possibly include such components as memory, input/output devices, and/or network interfaces, among others.

The term “continuous” as used herein, is a broad term encompassing its plain an ordinary meaning and, as used in reference to various types of activity (for example, scanning, monitoring, logging, and the like), includes without limitation substantially continuous activity and/or activity that may include periodic or intermittent pauses or breaks, but which accomplish the intended purposes described (for example, continuous scanning may include buffering and/or storage of data that is thereafter processed, for example, in batch and/or the like).

Any process descriptions, elements, or blocks in the flow diagrams described herein and/or depicted in the attached figures should be understood as potentially representing modules, segments, or portions of code which include one or more executable instructions for implementing specific logical functions or steps in the process. Alternate implementations

81

are included within the scope of the embodiments described herein in which elements or functions may be deleted, executed out of order from that shown or discussed, including substantially concurrently or in reverse order, depending on the functionality involved, as would be understood by those skilled in the art.

It should be emphasized that many variations and modifications may be made to the above-described embodiments, the elements of which are among other acceptable examples. All such modifications and variations are intended to be included herein within the scope of this disclosure. The foregoing description details certain embodiments of the invention. It will be appreciated, however, that no matter how detailed the foregoing appears in text, the invention may be practiced in many ways. As is also stated above, it should be noted that the use of particular terminology when describing certain features or aspects of the invention should not be taken to imply that the terminology is being re-defined herein to be restricted to including any specific characteristics of the features or aspects of the invention with which that terminology is associated. The scope of the invention should therefore be construed in accordance with the appended claims and any equivalents thereof.

What is claimed is:

1. A computer system for protecting a computer network from malware by providing for the efficient analysis of large amounts of malware-related data, the computer system comprising:

- one or more computer readable storage devices configured to store:
 - a plurality of computer executable instructions;
 - one or more software modules including computer executable instructions, the one or more software modules including a cluster engine module, a user interface engine module and a workflow engine module;
 - a plurality of data clustering strategies based on rules generated by the cluster engine module; and
 - a plurality of data cluster types, each data cluster type of the plurality of data cluster types associated with a data clustering strategy;
- one or more cluster data sources configured to store:
 - a plurality of data items including at least:
 - file data items, each file data item associated with at least one suspected malware file; and
 - malware-related data items associated with captured communications between an internal network and an external network, the malware-related data items including at least one of: external Internet Protocol addresses, external domains, external computing devices, internal Internet Protocol addresses, internal computing devices, users of particular computing devices, or organizational positions associated with users of particular computing devices; and
- one or more hardware computer processors in communication with the one or more computer readable storage devices and the one or more cluster data sources, and configured to execute the one or more software modules in order to cause the computer system to:
 - designate, by the cluster engine module, one or more seeds by:
 - accessing, from the one or more cluster data sources, the file data items;
 - calculating, for each file data item of the file data items, at least one of a hash of the file data item or a hash of an executed file data item, wherein the executed file

82

- data item was generated by an execution of the file data item in a sandboxed environment; and
- identifying one or more file data items based at least in part on comparing the at least one hash of the file data item or the executed file data item with a malware threat list of hashes, and designating each of the identified one or more file data items as a seed;
- for each of the file data items designated as a seed;
- select, by the cluster engine module, a particular data clustering strategy;
- identify, by the cluster engine module, one or more malware-related data items determined to be associated with the designated file data item seed based at least on the particular data clustering strategy, wherein the particular data clustering strategy performs at least one of querying the one or more cluster data sources or scanning network traffic to determine at least one of: external Internet Protocol addresses associated with the designated file data item seed, external domains associated with the designated file data item seed, external computing devices associated with the designated file data item seed, internal Internet Protocol addresses associated with the designated file data item seed, internal computing devices associated with the designated file data item seed users of particular computing devices associated with the designated file data item seed, or organizational positions associated with the determined users of particular computing devices;
- generate, by the cluster engine module, a data item cluster based at least on the designated file data item seed, wherein generating the data item cluster comprises:
 - adding the designated file data item seed to the data item cluster;
- identifying one or more of the network indicators that are associated with the seed;
- identifying one or more of the network-related data items associated with at least one of the identified one or more of the network indicators;
- adding, to the data item cluster, the identified one or more malware-related data items;
- identifying an additional one or more data items, including file data items and/or malware-related data items, associated with any data items of the data item cluster;
- adding, to the data item cluster, the additional one or more data items; and
- storing the one or more data item clusters,
- generating by the user interface engine module at least one human-readable conclusion associated with at least one generated data item cluster, wherein generating the at least one human-readable conclusion comprises:
 - determining a particular data cluster type from the plurality of cluster types based at least on the particular data clustering strategy;
 - identifying one or more human-readable templates comprising pre-generated text, wherein the human-readable templates are based at least on predefined associations between respective human-readable templates and data cluster types;
 - automatically analyzing the data item cluster to generate summary data according to rules, scoring algorithms, or other criteria; and
 - populating the identified one or more human-readable templates with data from the at least one generated

83

data item cluster or summary data of the at least one generated data item cluster;

cause presentation, by the user interface engine module, of the at least one generated data item cluster and the at least one human-readable conclusion including the populated pre-generated text data, in a user interface of a client computing device; and

generate, by the workflow engine and the user interface engine, an interactive workflow process to allow the user to perform at least one of: select new seeds, operate on existing seeds, generate new data clusters, or regenerate existing clusters.

2. The computer system of claim 1, wherein each of the data items of the data item cluster identify at least an internal computing device, a user of the internal computing device, and an organizational position associated with the user.

3. The computer system of claim 1, wherein the one or more hardware computer processors are further configured to execute the one or more software modules in order to cause the one or more hardware computer processors to:

scan communications between the internal network and the external network so as to identify additional malware-related data items; and

store the additional malware-related data items in the one or more cluster data sources.

4. The computer system of claim 3, wherein the communications are continuously scanned via a proxy.

5. The computer system of claim 1, wherein designation of the one or more seeds further comprises determining one or more network indicators associated with the identified one or more file data items, the one or more network indicators include at least one of an external Internet Protocol address or an external domain.

6. The computer system of claim 5, wherein the one or more of the network indicators that are associated with the seed comprise network indicators that are contacted by the at least one suspected malware file associated with the seed when the at least one suspected malware file is executed.

7. The computer system of claim 1, wherein designation of the one or more seeds further comprises determining whether or not a respective file data item has been marked by a human analyst as a seed.

8. The computer system of claim 7, wherein each of the file data items is processed by the computer system by at least:

initiating an analysis of the file data item including the at least one suspected malware file, wherein the analysis of the file data item generates a plurality of analysis information items including at least one of calculated hashes, file properties, academic analysis information, file execution information, or third-party analysis information;

associating the plurality of analysis information items with the file data item; and

generating a user interface including one or more user selectable portions presenting various of the analysis information items, the user interface usable by the human analyst to determine one or more characteristics of the file data item and to mark the file data item as a seed.

9. The computer system of claim 8, wherein the file data item is marked by a human analyst as a seed via a user interface of the computer system.

10. The computer system of claim 8, wherein initiating analysis of the file data item and generating the plurality of analysis information items comprises:

initiating an internal analysis of the file data item; and

initiating an external analysis of the file data item,

84

wherein the internal analysis includes analysis performed by the one or more hardware computer processors, and wherein the internal analysis includes at least one of calculation of an MD5 hash of the file data item, calculation of a SHA-1 hash of the file data item, or calculation of a size of the file data item, and

wherein the external analysis includes analysis performed by at least a second computer system, and wherein the external analysis includes execution of the file data item in a sandboxed environment and analysis of the file data item by a third-party malware analysis service.

11. The computer system of claim 1, wherein:

the one or more computer readable storage devices are further configured to store:

a plurality of data cluster analysis rules associated with a respective data clustering strategy, and

the one or more hardware computer processors are further configured to execute the one or more software modules in order to cause the one or more hardware computer processors to:

for each generated data item cluster:

access a plurality of data cluster analysis rules associated with the data clustering strategy;

analyze the data item cluster based on the accessed data cluster analysis rules; and

based at least on the analysis of the data item cluster;

determine an alert score for the data item cluster; and

cause presentation of the determined alert score in the user interface

generate one or more human-readable conclusions regarding the data item cluster.

12. The computer system of claim 11, wherein the alert score indicates a degree of correlation between characteristics of the data item cluster and the accessed data cluster analysis rules.

13. The computer system of claim 12, wherein a relatively higher alert score indicates a data cluster that is relatively more important for a human analyst to evaluate, and a relatively lower alert score indicated a data cluster that is relatively less important for the human analyst to evaluate.

14. The computer system of claim 12, wherein each alert score for respective data clusters is assigned to a category indicating a high degree of correlation, a medium degree of correlation, or a low degree of correlation.

15. The computer system of claim 14, wherein the high degree of correlation is associated with a first color, the medium degree of correlation is associated with a second color, and the low degree of correlation is associated with a third color.

16. The computer system of claim 11, wherein the one or more hardware computer processors are further configured to execute the one or more software modules in order to cause the computer system to:

for each generated data item cluster:

generate one or more human-readable conclusions regarding the data item cluster;

generate an alert, the alert comprising a the alert score, the one or more human-readable conclusions, the data items associated with the data item cluster, and metadata associated with the data items of the data item cluster.

17. The computer system of claim 16, wherein the one or more hardware computer processors are further configured to execute the one or more software modules in order to cause the computer system to:

generate a second user interface including a list of user-selectable alert indicators, an alert indicator being provided for each of the generated alerts, each of the alert

85

indicators providing a summary of information associated with respective generated alerts.

18. The computer system of claim 17, wherein the one or more hardware computer processors are further configured to execute the one or more software modules in order to cause the computer system to:

in response to a selection of an alert indicator by a human analyst;

generate an alert display, the alert display including at least an indication of the alert score and a list of the one or more human-readable conclusions.

19. The computer system of claim 1, wherein the at least human-readable conclusion further comprises a phrase or sentence including one or more indications of summary or aggregated data associated with a plurality of the data items of the at least one generated data item cluster.

* * * * *

86